The Fiscal Year 2007 Department of Defense

Internet Protocol Version 6

Test and Evaluation Report

September 2007

Assistant Secretary of Defense for Networks and Information Integration/
Department of Defense Chief Information Officer

UNCLASSIFIED

The Fiscal Year 2007 Department of Defense Internet Protocol Version 6 Test and Evaluation Report

This report is provided in response to Section 221 of the National Defense Authorization Act for Fiscal Year 2006 (Public Law 109-163). The report provides assessments of the test and evaluation results that the Department of Defense (DoD) Components have submitted to the DoD for the period July 2006 through June 2007 and integrates these assessments with the results previously reported by the DoD to Congress. The assessments follow the processes and methodologies of the test and evaluation strategy set forth in the Department of Defense Internet Protocol Version 6 Master Test Plan Version 2.0.

Approved by: John G. Grimes

John G. Grimes
Assistant Secretary of Defense for
Networks and Information Integration/
DoD Chief Information Officer

Dated: 14 September 2007

Approved by: Charles E McQueary

Dr. Charles E. McQueary
Director, Operational Test and Evaluation

Dated: 22 August 2007

Table of Contents

List of Tables

Executive Summary

This report provides a response to Section 221 of Public Law 109-163. It is based on field tests, exercises, demonstrations, experiments, simulations, and analyses conducted by Department of Defense (DoD) Components over the last five years, with emphasis on the most recent year (July 2006 through June 2007) test results. This report provides an update to the report submitted to Congress at the end of the last Fiscal Year (FY).

The DoD Internet Protocol (IP) Version 6 (IPv6) Transition Office (DITO) established a repository of IPv6 Test and Evaluation (T&E) reports provided by DoD Components in response to requests from the Assistant Secretary of Defense for Networks and Information Integration/ DoD Chief Information Officer (ASD(NII)/DoD CIO). The data contained in these reports have been evaluated with respect to the principal T&E objectives of the DoD IPv6 Master Test Plan Version 2.0 (MTP v2.0). The DoD Components [Army, Navy, Air Force, National Security Agency (NSA), and Defense Information Systems Agency (DISA)] have provided a total of 102 reports. For FY 2007, 44 reports were received, 19 reports for FY 2006, and 39 reports for FY 2005 (for T&E conducted FY 2003 through FY 2005).

The DoD Components reported a significant increase of IPv6 T&E activity during this reporting period, covering all 10 Joint Staff IPv6 operational criteria. Based on a cumulative analysis of all reports, one of the ten criteria, scalability (Criterion 6), has been fully demonstrated for transition to IPv6. Interoperability (Criterion 2) and performance (Criterion 3) are expected to be completely demonstrated in the upcoming year as well as elements of network transition techniques (Criterion 8). More T&E in operationally-realistic environments is needed to verify the demonstration of these criteria.

Although there was considerable T&E for security (Criterion 1) during this reporting period, commercial development and implementation of security devices/applications are still needed in demonstrating this criterion. Voice, data, and video integration (Criterion 4) and operation in low-bandwidth environment (Criterion 5) need technical guidelines, defined standards, and products (available for Criterion 4) to further demonstrate these criteria. Mobility (Criterion 7), network management (Criterion 9), and ad hoc networking (Criterion 10) lacked development and implementation, resulting in limited T&E.

The DoD Components are developing T&E plans for their specific Joint Staff IPv6 operational criteria and are following the guidance set forth in the DoD IPv6 MTP v2.0. The DoD is facilitating the sharing of IPv6 T&E results among DoD Components and other federal IPv6 working groups through DoD web portals.

The results presented in this report indicate that while considerable T&E is still required on operational networks, the maturity and stability of IPv6 technologies have made significant progress. T&E is required to support demonstration of the Joint Staff IPv6 operational criteria and DoD Approved Products List (APL) certification. The development and availability of critical, fully functional IPv6-capable products lag in some areas that could affect the DoD's schedule for IPv6 T&E and deployment.

1 Introduction

1.1 Purpose

The FY 2007 DoD IPv6 T&E Report is provided in response to Section 221 of Public Law 109-163. This report provides an assessment of IPv6 T&E activities carried out by the DoD Components with respect to the T&E objectives of the DoD IPv6 MTP v2.0. This report is also an input to the congressionally directed IPv6 certification by the Chairman of the Joint Chiefs of Staff.

1.2 Test and Evaluation Objectives

The DoD IPv6 T&E Report provides consolidated test results and assessments in support of the DoD transition to IPv6 and identifies what is completed and what T&E is still required. Assessment of the individual IPv6 T&E reports furnished by the DoD Components will address the progress in meeting the two objectives, as defined in the DoD IPv6 MTP v2.0:

- Demonstrate the functionality of IPv6 as delineated in the Joint Staff IPv6 operational criteria.
- Establish an APL of IPv6 products that have been certified to meet a set of DoD requirements for interoperability and Information Assurance (IA).

1.2.1 Demonstration of the Joint Staff IPv6 Operational Criteria

The Joint Staff enumerated 10 operational criteria that must be demonstrated before the DoD transitions its networks to IPv6. These criteria provide the top-level operational and technical capabilities necessary to verify that IPv6 fulfills the needs of the DoD. Each criterion was decomposed to provide two subordinate levels of measurable and verifiable functional elements that allow demonstration through T&E:

- Level 1 decomposition identifies capabilities required for each criterion.
- Level 2 decomposition identifies the specific technology, infrastructure, and/or functionality to demonstrate Level 1 decomposition.

The criteria and their associated Level 1 and Level 2 decomposition elements have been allocated among the Military Departments, NSA, and DISA for further decomposition and subsequent test coordination.

Additionally, Congress directed the Chairman of the Joint Chiefs of Staff to provide certification that conversion of DoD networks to IPv6 would “provide equivalent or better performance and capabilities than that which would be provided by any other combination of available technologies and protocols." The mapping of the DoD Components' IPv6 T&E results to the Joint Staff IPv6 operational criteria will support this certification.

1.2.2 Approved Products List

The DoD APL is a registry of information technology products which have been assessed by DoD entities and have passed DoD interoperability and information assurance (IA) requirements. Beginning in 2008, IPv6 capability will be assessed for all information technology products submitted for inclusion on the DoD APL. The addition of an information technology product to the DoD APL will occur only after the product meets DoD IPv6 certification requirements. Requirements for IPv6 interoperability certifications derive from the DoD Information Technology Standards Registry (DISR) IPv6 Standard Profiles for IPv6 Capable Products. The processes, procedures, and technical standards for the IA portion of testing are currently under development. Once developed, products will be tested for IA compliance. DoD Components shall purchase information technology products from the DoD APL.

DISA Joint Interoperability Test Command (JITC) is responsible for interoperability testing processes and procedures for products that are placed on the APL. IPv6-capable products are divided into seven categories: host, network appliance, router, applications, layer 3 (L3) switch, security device, and network server. A growing number of products are listed on the DoD APL, including one host, one network appliance, 17 routers, one web browser, one mail client, and one network server. DISA is responsible for developing processes, procedures, and technical standards for IPv6 IA testing. The DoD APL is located at: http://jitc.fhu.disa.mil/adv_ip/register/register.html.

1.3 Scope

The scope of analysis in this report is limited to T&E reports submitted by DoD Components in response to requests from the ASD(NII)/DoD CIO. The DoD received 44 reports (the most significant in terms of supporting the T&E objectives to date) from the Components (Army, Navy, Air Force, NSA, and DISA) during FY 2007, 19 reports for FY 2006, and 39 reports for FY 2005 (for T&E conducted FY 2003 through FY 2005). The evaluation team for this report was led by DISA (JITC), under the direction of ASD(NII)/DoD CIO and Director, Operational Test and Evaluation (DOT&E), and supported by DITO. This report provides the results of analyses for the 44 reports and integrates the analyses with the 58 previously submitted reports to provide a cumulative status for IPv6 T&E. This year's cumulative status is compared with the last two years to assess progress toward IPv6 transition.

1.4 FY 2005 and FY 2006 Reported Results and Recommendations

The FY 2005 DoD IPv6 T&E Report indicated that IPv6 technologies, as examined by the DoD Components, had progressed significantly toward the point of adoption and that some aspects of IPv6 appeared ready to deploy in a single network domain or enclave environment within operational networks. However, results and recommendations from that report indicated that additional effort was needed in the areas of security, performance, scalability, creation of a DoD APL, application porting or development, Quality of Service (QoS), transition mechanisms, and network management.

All areas identified in the FY 2005 T&E Report as needing additional T&E efforts, with the exception of scalability, were examined in the FY 2006 T&E Report as follows: Interoperability (Criterion 2) and network transition techniques (Criterion 8) had progressed sufficiently to allow use of the base protocol and the major transition mechanisms (dual stack and tunneling) to support broader testing in more operationally-realistic environments. FY 2007 T&E was expected to encompass security (Criterion 1); performance (Criterion 3); and voice, data, and video integration (Criterion 4). IPv6 capabilities for other criteria [low-bandwidth operation (Criterion 5), and tactical deployability and ad hoc networking (Criterion 10)] were found to be still too immature to support significantly expanded testing.

2 IPv6 Test and Evaluation Results

2.1 Overview

This section provides the overall status of DoD IPv6 T&E in support of the DoD's transition to IPv6 and summarizes IPv6 T&E results reported by DoD Components for the period July 2006 through June 2007. There were 44 T&E reports analyzed for the current reporting period. This was a significant increase in the number of reports from the previous year. Appendix D contains the summaries for each of these reports. Reports submitted for the current reporting period address the Joint Staff IPv6 operational criteria more clearly and generally present better T&E results than the previous year. T&E indicates that elements of three criteria (interoperability, scalability, and transition mechanisms) have been demonstrated with a high confidence factor. All reports used for this analysis can be found on the DoD Test and Evaluation Working Group (TEWG) portal: https://gesportal.dod.mil/sites/JITCIPv6/TEWG.

2.2 Cumulative Analysis Methodology

The cumulative status of each Joint Staff IPv6 operational criterion is based on analysis of all applicable tests conducted by DoD Components and is represented by a pie chart with slices colored red, yellow, or green. Each slice of a criterion's pie represents one Level 2 decomposition element for that criterion. The status color for each Level 2 element is based on analysis and evaluation of three factors as described in section 2.3 of this document. Underlying decomposition elements needing additional T&E are easily identified.

The color-coded rating scale for the Level 2 decomposition elements is as follows:

- Red - Limited progress has been made. A red slice indicates a Level 2 decomposition element that has had little or no T&E, or for which existing T&E results are inconclusive or unsatisfactory. Significant T&E and/or development is needed.

- Yellow - Significant progress has been made. A yellow slice indicates a Level 2 decomposition element that has had considerable T&E and for which multiple, independent T&E have provided substantially similar, positive results. But some combination of additional analysis, testing, or development is needed.

- Green - Successfully demonstrated. A green slice indicates a Level 2 decomposition element that has been successfully demonstrated. The evaluation type, relevance, and scope (considered with the number of tests) provide enough data to yield a high confidence factor.

The Cumulative Test and Evaluation Matrix (Table 2-1) presents the total number of T&E reports applicable to each criterion, categorized by evaluation method for the entire transition effort (counts for this reporting period are in parentheses). A cumulative pie chart through 2006 and 2007 is presented for each criterion representing the overall effort. An anticipated completion date to fully demonstrate the criterion is also provided. The cumulative pie charts provide the proportion of each criterion at each status level. A cumulative pie chart that is mostly red should be viewed as an alert that the demonstration of the underlying functional or technical elements is incomplete. A cumulative pie chart that is mostly yellow means that most underlying elements have had considerable progress. A cumulative pie chart that is all green indicates that all underlying elements for that criterion were fully tested and the criterion has been satisfactorily demonstrated.

Table 2-1 Cumulative Test and Evaluation Matrix

Joint Staff IPv6 Operational Criteria		Test Methods							Cumulative Status Thru		Expected Completion Date
		Engineering Analyses	Modeling & Simulation	Experiments	Demonstrations	Pilots	Exercises	Field Tests	2006	2007	
1	Demonstrate security of unclassified network operations, classified network operations, black backbone operations, integration of High Assurance IP Encryptors (HAIPE), integration of IP security (IPSec), and integration with firewalls and intrusion detection systems	20 (11)	1	15 (8)	8 (4)	2 (2)	11 (4)		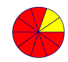	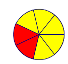	1QFY 2009
2	Demonstrate end-to-end interoperability in a mixed IPv4 and IPv6 environment	11 (6)	2	13 (1)	7 (2)	1 (1)	17 (4)	1	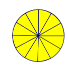	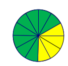	3QFY 2008
3	Demonstrate equivalent to, or better performance than, IPv4 based networks	2	2	4 (1)	4 (4)		8 (1)		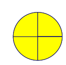	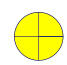	1QFY 2008
4	Demonstrate voice, data, and video integration	6 (2)		2			4 (1)	1	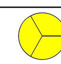	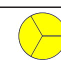	4QFY 2008
5	Demonstrate effective operation in low-bandwidth environment	2	2		1 (1)		5 (2)		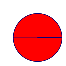	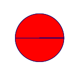	2QFY 2009
6	Demonstrate scalability of IPv6 networks	1 (1)		1 (1)	1	1 (1)			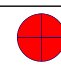	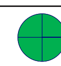	1QFY 2008
7	Demonstrate support for mobile terminals (voice, data and video)	5 (2)	1	1	1 (1)		7	1	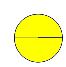	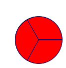	2QFY 2009
8	Demonstrate transition techniques	16 (8)	4 (1)	20 (10)	8 (5)	2 (2)	20 (8)		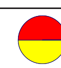	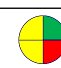	4QFY 2010
9	Demonstrate ability to provide network management of networks	3 (1)		6	4 (1)				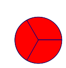	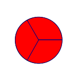	3QFY 2008
10	Demonstrate tactical deployability and ad hoc networking	7 (4)	1	2 (1)	1 (1)			1	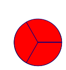	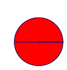	2QFY 2010

Key:

- Criterion has been successfully demonstrated.
- Significant progress has been made on this criterion.
- Limited progress has been made on this criterion.

The pie charts for criteria 1, 7, 8, and 10 differ from 2006 due to the change of Level 2 decomposition items. Refer to each criterion in Section 2.3 for more detail.

QFY Quarter Fiscal Year   Total Events (Current Fiscal Year Events)

2.3 Impact of FY 2007 Test and Evaluation Reports on Demonstration of Joint Staff IPv6 Operational Criteria

This section provides the evaluation of each Joint Staff IPv6 operational criterion at the lowest level of the decomposed functional or technical elements. Three qualitative factors were used to determine the extent to which an individual report contributed to the satisfaction of an element: applicability to the Joint Staff IPv6 operational criteria, qualitative merit based on evaluation type, and scope of each T&E event.

First, each T&E event was evaluated for applicability or relevance to each Joint Staff IPv6 operational criterion and for the degree of relevance that each event contributed to determining the Level 2 status. Next, the type of evaluation was considered and the event results were weighted accordingly. Evaluation types listed in descending qualitative order are: field test, exercise, pilot, demonstration, experiment, modeling and simulation, and engineering analysis. Finally, the scope of each T&E event was considered. In determining status, T&E events that only confirm previous results were allocated less weight than those that cover previously untested areas.

The color-coded rating scale used in the individual criterion's decomposition table is as follows:

⊗ Red - Limited progress has been made. More T&E and/or development is needed to allow the decomposition item to be certified as having been demonstrated or T&E to date has not demonstrated satisfactory results.

⊕ Yellow - Significant progress has been made. Some portions of the decomposition item have not been successfully demonstrated or confidence in previous T&E results was low. Additional T&E and/or development are needed to allow the decomposition item to be certified as having been demonstrated.

✓ Green - The decomposition item has been successfully demonstrated. T&E has provided enough data to assure the decomposition item was demonstrated with a high confidence factor.

Subsections follow for each criterion. Each subsection provides the Level 1 and Level 2 decomposition status of each criterion through 2006 and 2007. Specific T&E observations related to that criterion for 2007 follow each table.

2.3.1 Criterion 1: Demonstrate security of unclassified network operations, classified network operations, black backbone operations, integration of HAIPE, integration of IPSec, and integration with firewalls and intrusion detection systems

Table 2-2 Joint Staff IPv6 Operational Criterion 1 Status

Level 1 Decomposition (Capabilities to be demonstrated)	Cumulative Status Thru		Level 2 Decomposition (Specific technology/infrastructure/ functionality to be demonstrated)	Cumulative Status Thru	
	2006	2007		2006	2007
1.1 Ensure that information is not disclosed to unauthorized persons, processes, or devices.	⊗	⊕	**1.1.1** Verify the implementation of IPSec with Encapsulating Security Protocol (ESP) in IPv6 hosts and routers. Verify integration with Public Key Infrastructure (PKI).	⊗	⊕[1]
1.2 Ensure information received is the same as that which was sent (protect against unauthorized modification or destruction of information).	⊗	⊕	**1.2.1** Verify implementation of Authentication Header (AH) in IPv6 hosts and routers. Verify integration with PKI.	⊗	⊕[1]
1.3 Ensure Authentication, Authorization, and Accounting (AAA) of persons and processes.	⊗	⊗	**1.3.1** Verify the implementation of a AAA server is able to ensure the Authentication, Authorization, and Accounting of persons, machines, and processes over an IPv6 network.	⊗	⊗
1.4 Ensure availability and mitigate denial of services (timely, reliable access to data, and information services for authorized users).	⊕	⊕	**1.4.1** Verify protection of the IPv6 stack of Hosts and Network Devices from intruders. (Note: Included in this are vulnerabilities that arise from errors in protocol specification or implementation or the associated device firmware).	⊕	⊕
			1.4.2 Demonstrate IPv6 traffic filtering capabilities of routers and firewalls according to security policies.	⊕	⊕
1.5 Ensure IPv6 traffic is interoperable with firewalls and Intrusion Detection Systems (IDS).	⊗	⊕	**1.5.1** Evaluate Firewalls and IDS functions that can be applied to IPv6 traffic. Evaluate Firewalls and IDS functions that can be applied to tunneled IPv6 traffic.	⊗	⊕
1.6 Ensure IPv6 traffic is interoperable with HAIPE devices.	⊗	⊗	**1.6.1** Evaluate HAIPE v3's ability to encrypt/decrypt IPv6 packets.	⊗	⊗

[1] Level 2 Decomposition 1.1.1 and 1.2.1 were modified in 2007, changing the number of Level 2 elements.

2007 T&E Observations

- Few products fully support IPv6 IPSec; however, vendors have implemented IPSec on intermediate systems (i.e., routers and L3 switches).
  (Test Report D.4; Decomposition 1.1.1)

- All products tested during this reporting period for the DoD APL support AH which is the part of IPSec that is defined in Request For Comment (RFC) 4302.
  (Test Reports D.17, 23, 24; Decomposition 1.2.1)

- Implementing ESP within hosts and routers was successfully demonstrated; however, there was insufficient T&E of Internet Key Exchange (IKE).
  (Test Reports D.17, 23; Decomposition 1.1.1)

- In FY 2006, vulnerability testing against IPv6 yielded 84 vulnerabilities among five operating systems [OS(s)]. During mitigation testing, 85% of the vulnerabilities passed testing after installation of vendor patches, implementation of Secure Technical Implementation Guides (STIGs) and custom configurations. Only 13 vulnerabilities remained.
  (Test Reports D.40-44; Decomposition 1.4.1)

- Tested vendor OSs showed varying behavior in response to router advertisement attacks, which could lead to denial of services.
  (Test Report D.8; Decomposition 1.4.1, 1.4.2)

- Firewall T&E produced mixed results. One product was found to support IPv6 and was able to simultaneously provide stateful inspection of both IPv4 and IPv6 data streams with little or no negative performance impact. However, testing by NSA revealed another firewall did not provide adequate IPv6 functionality.
  (Test Report D.18, Unpublished NSA Report; Decomposition 1.5.1)

- Although the HAIPE v3 specifications include IPv6 requirements, none were tested because IPv6-capable HAIPE devices are still under development.
  (Test Report D.5; Decomposition 1.6.1)

2.3.2 Criterion 2: Demonstrate end-to-end interoperability in a mixed IPv4 and IPv6 environment

Table 2-3 Joint Staff IPv6 Operational Criterion 2 Status

Level 1 Decomposition (Capabilities to be demonstrated)	Cumulative Status Thru 2006	Cumulative Status Thru 2007	Level 2 Decomposition (Specific technology/infrastructure/ functionality to be demonstrated)	Cumulative Status Thru 2006	Cumulative Status Thru 2007
2.1 Demonstrate IPv4 application to IPv4 application over a mixed IPv4 and IPv6 network.	⊕	✓	**2.1.1** Demonstrate core service interoperability: Domain Name System (DNS), directory services, File Transfer Protocol (FTP), email, web services, Network Time Protocol (NTP), and PKI.	⊕	✓
			2.1.2 Demonstrate network core application interoperability: Voice over IP (VoIP) and video over IP.	⊕	✓
			2.1.3 Demonstrate Commercial Off The Shelf (COTS) application interoperability (transaction, database access, and web services).	⊕	✓
			2.1.4 Demonstrate Government Off The Shelf (GOTS) applications/systems interoperability.	⊕	✓
2.2 Demonstrate IPv6 application to IPv4 application over a mixed IPv4 and IPv6 network.	⊕	⊕	**2.2.1** Demonstrate core service interoperability: DNS, Directory, FTP, email, web services, NTP, and PKI.	⊕	⊕
			2.2.2 Demonstrate network core application interoperability: VoIP and video over IP.	⊕	⊕
			2.2.3 Demonstrate COTS application interoperability (transaction, database access, and web services).	⊕	⊕
			2.2.4 Demonstrate GOTS application/system interoperability.	⊕	⊕
2.3 Demonstrate IPv6 application to IPv6 application over a mixed IPv4 and IPv6 network.	⊕	✓	**2.3.1** Demonstrate core service interoperability: DNS, Directory, FTP, email, web services, NTP, and PKI.	⊕	✓
			2.3.2 Demonstrate network core application interoperability: VoIP and video over IP.	⊕	✓
			2.3.3 Demonstrate COTS application interoperability (transaction, database access, and web services).	⊕	✓
			2.3.4 Demonstrate GOTS application/system interoperability.	⊕	✓

2007 T&E Observations

- Core services DNS, FTP, email, VoIP, and video over IP successfully interoperated in many mixed IPv4 and IPv6 environments.
(Test Reports D.1, 2, 4, 16, 19, 23; Decomposition 2.2.1, 2.2.2, 2.3.1, 2.3.2)

- The Warfighter Information Network-Tactical (WIN-T) system test demonstrated VoIP, email exchange, and FTP sessions with over 99% of voice calls and 100% of data exchanges completing successfully in a mixed IPv4 and IPv6 environment.
(Test Report D.2; Decomposition 2.2.1, 2.2.2, 2.3.1, 2.3.2)

- In one test event, all connection-oriented Transmission Control Protocol (TCP) scripts successfully demonstrated the transport and delivery for all protocols and traffic types tested. Connectionless User Datagram Protocol (UDP) scripts exhibited a 99%+ success rate while running concurrently with connection-oriented scripts.
(Test Report D.16; Decomposition 2.2.1, 2.2.2, 2.3.1, 2.3.2)

- Testing indicated the following protocols and applications to be interoperable:
(Test Report D.19; Decomposition 2.2.1, 2.2.2, 2.3.1, 2.3.2)

    - FTP (Get/Put)
    - Hypertext Transfer Protocol (HTTP)
    - HTTP Secure (HTTPS)
    - Post Office Protocol version 3 (POP3)
    - Simple Mail Transfer Protocol (SMTP)
    - Simple Network Management Protocol (SNMP)
    - Lightweight Directory Access Protocol (LDAP)
    - Session Initiation Protocol (SIP)
    - DNS
    - G.711u (VoIP)
    - IP Television (IPTV) – Video
    - IPTV – Audio

- There has been no T&E of IPv6 GOTS user-level applications; there is little demand for these applications at this time.
(General Observation; Decomposition 2.1.3, 2.2.3, 2.3.3)

- Interoperability could not be demonstrated for NTP, Dynamic Host Configuration Protocol (DHCP), Resource Reservation Protocol (RSVP), and PKI due to lack of implementation or maturity in IPv6.
(Test Reports D.16, 19; Decomposition 2.2.1, 2.2.2, 2.3.1, 2.3.2)

- The implementation of DHCP version 6 (DHCPv6) remains a T&E issue from last year due to the need for continuing development of protocols and vendor products.
(Test Reports D.16, 19; Decomposition 2.1.1, 2.2.1, 2.3.1)

2.3.3 Criterion 3: Demonstrate equivalent to, or better performance than, IPv4 based networks

Table 2-4 Joint Staff IPv6 Operational Criterion 3 Status

Level 1 Decomposition (Capabilities to be demonstrated)	Cumulative Status Thru		Level 2 Decomposition (Specific technology/infrastructure/ functionality to be demonstrated)	Cumulative Status Thru	
	2006	2007		2006	2007
3.1 Demonstrate IPv6 throughput equivalent to or better than IPv4.	⊕	⊕	**3.1.1** Same as Level 1.	⊕	⊕
3.2 Demonstrate IPv6 latency equivalent to or better than IPv4.	⊕	⊕	**3.2.1** Same as Level 1.	⊕	⊕
3.3 Demonstrate IPv6 packet loss equivalent to or better than IPv4.	⊕	⊕	**3.3.1** Same as Level 1.	⊕	⊕
3.4 Demonstrate IPv6 service availability equivalent to or better than IPv4.	⊕	⊕	**3.4.1** Same as Level 1.	⊕	⊕

2007 T&E Observations

- The performance of native IPv6 traffic and dual-stack traffic over the Global Broadcasting System (GBS) IPv6 pilot architecture has proven to be more efficient than the current IPv4 architecture.
  (Test Report D.11; Decomposition 3.1.1, 3.2.1, 3.3.1)

- Performance testing during the IPv6 Low-Bandwidth Test showed IPv6 performance equivalent to IPv4 on the specified routers and multiplexers.
  (Test Report D.19; Decomposition 3.1.1, 3.2.1, 3.3.1)

- T&E demonstrated equivalent performance exists between the frame size, throughput, and latency values for IPv4 and IPv6.
  (Test Report D.38; Decomposition 3.1.1, 3.2.1)

- In the “2006 Ethernet Switch Comparison Report,” equipment from all six vendors demonstrated performance at or near line rate when processing IPv6 traffic, which is similar to IPv4 test results.
  (Test Report D.34; Decomposition 3.1.1)

- The IPv6 and IPv4 network performance characteristics of throughput and latency were virtually identical in low-bandwidth environments of 8 to 512 Kilobits per second (Kbps). Differences were generally less than 1%.
  (Test Report D.19; Decomposition 3.1.1, 3.2.1)

- All planned T&E to demonstrate this criterion will be completed by the end of FY 2007. (General Observation; Decomposition 3.1.1, 3.2.1, 3.3.1)

2.3.4 Criterion 4: Demonstrate voice, data, and video integration

Table 2-5 Joint Staff IPv6 Operational Criterion 4 Status

<table>
<tr><th rowspan="2">Level 1
Decomposition
(Capabilities to be demonstrated)</th><th colspan="2">Cumulative Status Thru</th><th rowspan="2">Level 2
Decomposition
(Specific technology/infrastructure/ functionality to be demonstrated)</th><th colspan="2">Cumulative Status Thru</th></tr>
<tr><th>2006</th><th>2007</th><th>2006</th><th>2007</th></tr>
<tr><td rowspan="3">4.1 Demonstrate simultaneous voice, data, and video (or any combination thereof) over shared IPv6 networks.</td><td rowspan="3">⊕</td><td rowspan="3">⊕</td><td>4.1.1 Demonstrate Quality of Service (QoS) capabilities of IPv6 networks using Differentiated Services (DiffServ) and Resource Reservation Protocol (RSVP).</td><td>⊕</td><td>⊕</td></tr>
<tr><td>4.1.2 Demonstrate transport control capabilities of IPv6 networks using Real Time Protocol (RTP).</td><td>⊕</td><td>⊕</td></tr>
<tr><td>4.1.3 Demonstrate session signaling capabilities of IPv6 networks using the Session Initiation Protocol (SIP).</td><td>⊕</td><td>⊕</td></tr>
</table>

2007 T&E Observations

- DiffServ was successfully demonstrated using multiple data streams with various designated levels of service.
  (Test Report D.26; Decomposition 4.1.1)

- T&E of several layer 3 switches demonstrated QoS prioritization capabilities of IPv6 equivalent to IPv4.
  (Test Report D.34; Decomposition 4.1.1)

- Applications for RTP and SIP, including the closely associated Assured Services-SIP (AS-SIP), are in development, resulting in limited demonstration of this criterion.
  (General Observation; Decomposition 4.1.2, 4.1.3)

2.3.5 Criterion 5: Demonstrate effective operation in low-bandwidth environment

Table 2-6 Joint Staff IPv6 Operational Criterion 5 Status

Level 1 Decomposition (Capabilities to be demonstrated)	Cumulative Status Thru 2006	Cumulative Status Thru 2007	Level 2 Decomposition (Specific technology/infrastructure/ functionality to be demonstrated)	Cumulative Status Thru 2006	Cumulative Status Thru 2007
5.1 Demonstrate ability to establish and maintain applications in low-bandwidth IPv6 environments.	⊗	⊗	**5.1.1** Demonstrate ability to establish and maintain applications (voice, data, video) in low-bandwidth IPv6 environments.	⊗	⊗
			5.1.2 Demonstrate ability to maintain network operations (i.e., Network Management, DNS, Dynamic DNS, and Security) in low-bandwidth IPv6 environments.	⊗	⊗

2007 T&E Observations

- Testing in specific low-bandwidth windows (8, 16, 32, 64, 128, 256, and 512 Kbps) revealed an average variance of 0.7% packet latency between IPv4 and IPv6 packet transmission rates.
  (Test Report D.19; Decomposition 5.1.1)

- During QoS testing in a limited-bandwidth environment (256 Kbps), voice, data, and video streams demonstrated the expected level of service.
  (Test Report D.26; Decomposition 5.1.2)

- Multiple simulated VoIP calls were made across a 512 Kbps link (carrying other data and video traffic) with a 99.93% success rate.
  (Test Report D.2; Decomposition 5.1.1)

2.3.6 Criterion 6: Demonstrate scalability of IPv6 networks

Table 2-7 Joint Staff IPv6 Operational Criterion 6 Status

Level 1 Decomposition (Capabilities to be demonstrated)	Cumulative Status Thru 2006	Cumulative Status Thru 2007	Level 2 Decomposition (Specific technology/infrastructure/ functionality to be demonstrated)	Cumulative Status Thru 2006	Cumulative Status Thru 2007
6.1 Demonstrate ability to add more network resources, services and users without negatively impacting existing users.	⊗	✓	**6.1.1** Demonstrate the ability to build IPv6 networks comparable in size to existing IPv4 networks, with equal or better performance.	⊗	✓
			6.1.2 Demonstrate the ability to populate IPv6 subnets with network elements in comparable numbers to existing IPv4 subnets, with equal or better performance.	⊗	✓
			6.1.3 Demonstrate the ability to create IPv6 multicast sessions whose sizes are comparable to existing IPv4 multicast sessions, with equal or better performance.	⊗	✓
			6.1.4 Demonstrate the ability to create IPv6 core services (DNS, Directory, FTP, email, Web, NTP, PKI) where the number of users are comparable to existing IPv4 core services, with equal or better performance.	⊗	✓

2007 T&E Observations

- An IPv6 network demonstrated the capability of being scaled using the production Research, Development, Test, and Evaluation (RDT&E) Defense Research and Engineering Network (DREN) IPv6 pilot network.
  (Test Reports D.3, 4; Decomposition 6.1.1)

- T&E specifically designed to address this criterion showed that scaling networks, subnets, and multicast sessions did not degrade resources for IPv6 relative to IPv4; dual stacking, however, does put additional stress on memory resources but this is to be expected.
  (Test Report D.30; Decomposition 6.1.1, 6.1.2, 6.1.3)

- T&E showed that in a mixed IPv4/IPv6 environment with varying frame sizes, Computer Processor Unit (CPU) performance was not degraded by increased network traffic.
  (Test Report D.30; Decomposition 6.1.1)

- T&E indicates that enabling basic IP services (HTTP and FTP) for IPv4 and IPv6 scales equally well for both protocols.
  (Test Report D.30; Decomposition 6.1.4)

- All planned IPv6 scalability T&E has been completed.
(Test Report D.30; Decomposition 6.1.1, 6.1.2, 6.1.3, 6.1.4)

2.3.7 Criterion 7: Demonstrate support for mobile terminals (voice, data, and video)

Table 2-8 Joint Staff IPv6 Operational Criterion 7 Status

Level 1 Decomposition (Capabilities to be demonstrated)	**Cumulative Status Thru**		**Level 2 Decomposition (Specific technology/infrastructure/ functionality to be demonstrated)**	**Cumulative Status Thru**	
	2006	**2007**		**2006**	**2007**
7.1 Demonstrate ability to establish and maintain IPv6 applications (voice, data, video) on the move.	⊕	⊗	**7.1.1** Demonstrate ability to initiate and maintain voice, data, or video applications using mobile terminals.	⊕	⊗
			7.1.2 Demonstrate ability to maintain network operations of mobile terminals (i.e., Network Management, DNS, Dynamic DNS, and Security).	⊕	⊗
			7.1.3 Demonstrate the ability to maintain connectivity of Mobile Nodes (MN) while On-The-Move (OTM) and network management of MN while OTM.	**N/A**[2]	⊗

2007 T&E Observations

- The Interoperable Networks for Secure Communications (INSC) project demonstrated the ability to maintain existing voice communications using Mobile IPv6 (MIPv6) in a coalition environment.
(Test Report D.13; Decomposition 7.1.1)

- The INSC project also demonstrated Home Agent (HA) autoconfiguration for MIPv6 using a unique solution developed by the project that may be superior to existing Internet Engineering Task Force solutions.
(Test Report D.13; Decomposition 7.1.1)

- INSC test data indicated 3 to 4.16 seconds of overall handoff latency with and without Route Optimization (RO). This is due to additional processing required by MIPv6. Additional test data indicated that, even with two handoffs per minute, the TCP throughput between the corresponding node and MN provided by MIPv6 with RO was approximately 5.5 Mbps (Megabits per second) compared to 6 Mbps. The fact that connectivity was maintained without user intervention is a major improvement in MIPv6.
(Test Report D.13; Decomposition 7.1.3)

[2] Level 2 Decomposition 7.1.3 was added in 2007; therefore, cumulative status for 2006 is Not/Applicable (N/A).

- As a result of modifying the decomposition from 2006, the cumulative status and the status of the individual decomposition elements changed from yellow to red for this reporting period. The decomposition modifications were based on additional analysis by the Army and the logical association of an additional Level 2 decomposition element from Criterion 10. A comprehensive test plan for mobility is being developed to address any setbacks of this criterion.
(General Observation)

2.3.8 Criterion 8: Demonstrate transition techniques

Table 2-9 Joint Staff IPv6 Operational Criterion 8 Status

Level 1 Decomposition (Capabilities to be demonstrated)	Cumulative Status Thru		Level 2 Decomposition (Specific technology/infrastructure/ functionality to be demonstrated)	Cumulative Status Thru	
	2006	2007		2006	2007
8.1 Demonstrate DoD recommended network transition techniques.	⊕	⊕	**8.1.1** Demonstrate the interoperability of IPv4 and IPv6 network transition techniques: • Dual stack everywhere in an autonomous system • Configured tunnels • Tunnel Broker.	⊕	✓
			8.1.2 Demonstrate the performance of IPv4 and IPv6 network transition techniques: • Dual stack everywhere in an autonomous system • Configured tunnels • Tunnel Broker.	**N/A**[3]	⊕
			8.1.3 Demonstrate the security of IPv4 and IPv6 network transition techniques: • Dual stack everywhere in an autonomous system • Configured tunnels • Tunnel Broker.	**N/A**[3]	⊕
8.2 Demonstrate DoD recommended application transition techniques.	⊗	⊗	**8.2.1** Demonstrate the interoperability of the IPv4 and IPv6 application transition techniques: • Stateless IP/Internet Control Message Protocol Translation (SIIT) • Bump in the Application Program Interface (BIA) • Bump in the Stack (BIS).	⊗	⊗

[3] Level 2 Decomposition 8.1.2 and 8.1.3 were added in 2007; therefore, cumulative status for 2006 is N/A.

2007 T&E Observations

- Limited T&E of translators, transformers, and tunnel broker transition mechanisms was conducted. These mechanisms performed well in both laboratory and tactical satellite environments.
  (Test Report D.6; Decomposition 8.1.1, 8.1.2)

- In this reporting period, dual IP stacks exhibited stable coexistence and provided exceptional flexibility with acceptable impacts.
  (General Observation; Decomposition 8.1.1, 8.1.2)

- T&E results indicate that to conduct a successful IPv6 pilot, every affected device in the system should be dual stack.
  (Test Report D.37; Decomposition 8.1.1, 8.1.2)

- For dual-stack traffic, IPv6 packets arrived before IPv4 packets. This illustrated the efficiency of routers and network devices within the pilot network to process and forward IPv6 traffic. This was attributed to the streamlined design of the IPv6 header.
  (Test Report D.11; Decomposition 8.1.2)

- A report revealed that tunnels could degrade performance especially when processing is done in software, so processing in Application Specific Integrated Circuits (ASICs) is preferred. Results from this report also showed IPSec over IPv6 tunnels had little impact on performance.
  (Test Report D.27; Decomposition 8.1.2, 8.1.3)

- Of the five recommended transition mechanisms defined last year, dual stack, manual configured tunnel, and automatic tunneling were the most commonly tested.
  (General Observation; Decomposition 8.1.1)

2.3.9 Criterion 9: Demonstrate ability to provide network management of networks

Table 2-10 Joint Staff IPv6 Operational Criterion 9 Status

Level 1 Decomposition (Capabilities to be demonstrated)	Cumulative Status Thru		Level 2 Decomposition (Specific technology/infrastructure/ functionality to be demonstrated)	Cumulative Status Thru	
	2006	2007		2006	2007
9.1 Demonstrate ability to monitor, configure, and account for IPv6 network resources.	⊗	⊗	**9.1.1** Demonstrate that IPv6 devices can be monitored by Network Management Systems (NMS) commonly used by the DoD.	⊗	⊗
			9.1.2 Demonstrate that NMS commonly used by the DoD can configure IPv6 devices.	⊗	⊗
			9.1.3 Demonstrate that IPv6 devices can be accounted by NMS commonly used by the DoD.	⊗	⊗

2007 T&E Observations

- Few software tools currently support an IPv6-only mode of network management. Most tools currently require an IPv4 interface. However, current IPv4 tools can be used for management of IPv6 resources in mixed IPv4/IPv6 environments.
  (General Observation; Decomposition 9.1.1)

- Ethernet switches tested during this reporting period are still deficient in IPv6 management capability.
  (Test Report D.34; Decomposition 9.1.1)

- Given the status of IPv6 network management tools, significant T&E could not be conducted during this reporting period.
  (General Observation)

2.3.10 Criterion 10: Demonstrate tactical deployability and ad hoc networking

Table 2-11 Joint Staff IPv6 Operational Criterion 10 Status

Level 1 Decomposition (Capabilities to be demonstrated)	**Cumulative Status Thru**		**Level 2 Decomposition (Specific technology/infrastructure/ functionality to be demonstrated)**	**Cumulative Status Thru**	
	2006	**2007**		**2006**	**2007**
10.1 Demonstrate ability to move IPv6 networks as a whole, without reconfiguration.	⊗	⊗	**10.1.1** Demonstrate the ability to move networks to other locations while maintaining connectivity via the original IPv6 addresses, using Network Mobility (NEMO).	⊗	⊗[4]
10.2 Demonstrate ability to support IPv6 networking without fixed router infrastructure.	⊗	⊗	**10.2.1** Demonstrate ability of IPv6 hosts to forward packets from peers, while on the move, using Mobile Ad hoc Networks (MANET) routing protocols.	⊗	⊗

2007 T&E Observations

- During the Joint User Interoperability Communications Exercise (JUICE), using NEMO, networks operated in an On-The-Move (OTM) capacity with minimal user intervention. (Test Report D.36; Decomposition 10.1.1)

- Research continues into MANET solutions for highly mobile users at the tactical edge network where fixed infrastructure is not available. Autoconfiguration technologies were investigated for utilization with mobile networking with stateless autoconfiguration showing potential to simplify mobile scenarios.
  (Test Report D.20; Decomposition 10.2.1)

- One report noted the clear usefulness of mobility technologies in military applications adding that some show relative maturity such as MANET Optimized Link State Routing. Others are still experimental such as MANET-Open Shortest Path First (OSPF). MANET multicasting showed improved results with Simple Multicast Forwarding supporting the demonstration of streaming video, VoIP, and chat. NEMO remains mostly experimental.
  (Test Report D.13; Decomposition 10.2.1)

[4] Level 2 Decomposition for Criterion 10 was modified in 2007, changing the number of Level 2 elements.

3 Conclusions

The following conclusions are based upon reviewing and integrating the results of the 44 FY 2007 T&E reports. The DoD made significant progress in IPv6 T&E during this reporting period. However, further development of IPv6 protocols and/or T&E are required on most of the Joint Staff IPv6 operational criteria. The conclusions are summarized according to the Joint Staff IPv6 operational criteria.

Criterion 1: Demonstrate security of unclassified network operations, classified network operations, black backbone operations, integration of HAIPE, integration of IPSec, and integration with firewalls and intrusion detection systems.

- HAIPE integration with IPv6 continues to be a major concern. Although the technical specifications for HAIPE v3 were released on August of 2006, as of June 2007 there were still no commercially available implementations to test.
- For the portion of IPv6 IPSec (AH) that has been implemented by vendors, T&E results indicate favorable compliance. However, for the other portions of IPSec (i.e. ESP and IKE) T&E was limited and remains a primary concern for transition. Complete implementation of IPSec is not expected for some time.
- Serious deficiencies exist in IPv6 functionality needed to support PKI (directory services, DNS, key management, administrative support, and vendor expertise).
- NSA initiated OS vulnerability assessment and mitigation testing. The latter demonstrated the ability to protect host workstations on IPv6 networks. Additional vulnerability assessments are required in this area. Similar efforts are necessary for routers, switches, and network security devices (firewalls, IDS, etc.).
- Some vulnerability assessment and mitigation tools have been evaluated for IPv6. Commercial development and T&E for security products (e.g., certification tools, firewalls, IDS, and IPS) is essential.

Criterion 2: Demonstrate end-to-end interoperability in a mixed IPv4 and IPv6 environment.

- IPv4 and IPv6 can coexist without adverse impact on network operations.
- T&E this reporting period demonstrated sufficient interoperability of network devices, services, and applications; however, some features such as DHCP lack maturity and vendor offerings.
- There has been no T&E of IPv6 GOTS user-level applications; it appears there is little demand for these applications at this time.

Criterion 3: Demonstrate equivalent to, or better performance than, IPv4 based networks.

- Performance T&E produced results that support IPv6 parity with IPv4 based networks.
- The lack of IPv6-capable satellite IP modems and accelerators prevents assessment of satellite links for this criterion.
- The current state of IPv6 used in tactical networks requires T&E before accurate performance comparisons with IPv4 can be made.
- Completion of service availability testing and end-to-end network performance testing using "end-user experience" metrics is required and is scheduled to be completed by the end of FY 2007.

Criterion 4: Demonstrate voice, data, and video integration.

- DiffServ for IPv6 provided QoS capability (prioritizing packets) for different classes of voice, video, and data traffic.
- Further development, T&E, and technical guidelines for integrating QoS methodologies are required to adequately demonstrate this criterion.
- DoD QoS requirements and policies are needed to guide IPv6 T&E efforts.
- T&E is required of IPv6 applications and products using RTP, SIP, and specifically AS-SIP with the execution of RSVP.

Criterion 5: Demonstrate effective operation in low-bandwidth environment.

- Limited T&E has shown applications (voice, data, and video) can be established and maintained in low-bandwidth environments of 8 to 512 Kbps.
- Further IPv6 T&E in bandwidth-constrained, operationally-realistic tactical environments is required to fully demonstrate this criterion.

Criterion 6: Demonstrate scalability of IPv6 networks.

- T&E has shown that IPv6 scales equivalently to IPv4. Networks, subnets, multicast sessions, and network services on a commercial platform with various packet sizes and mixed ratios of IPv4 to IPv6 traffic were evaluated.
- Routers, L3 switches, security appliances and servers may require an upgrade (e.g., memory resources and CPU) to provide for dual-stack capabilities.
- All planned T&E to support demonstration of this criterion is complete.

Criterion 7: Demonstrate support for mobile terminals (voice, data, and video).

- Increased capability in mobile node technology has been demonstrated. Routers were able to incorporate Home Agent (HA) functionality and maintain connectivity at the halt.

- Lack of development and implementation of mobile applications by industry limited T&E for this criterion.

Criterion 8: Demonstrate transition techniques.

- The interoperability and functionality of the following IPv6 network transition mechanisms have been successfully demonstrated: dual stack, configured tunnels, and tunnel broker.

- Dual stacking appears to create the most flexible strategy for the coexistence of IPv6 with IPv4 and is sufficiently stable to allow deployment of mixed networks.

- The overall approach for transitioning GOTS applications has not been determined. The strategy is to transition applications following the change of DoD core networks. However, the use of dual stack may obviate the need for transitioning legacy applications.

- The network environment and mission requirements must be considered in selecting a transition mechanism. Not all mechanisms are expected to perform equally in all circumstances, and regardless of performance, may have certain advantages depending on the mission objectives.

Criterion 9: Demonstrate ability to provide network management of networks.

- IPv6 network management (using IPv4 management tools) has been implemented to a limited extent. Further commercial development of native IPv6 network management tools and T&E is required to demonstrate this criterion.

- Network management requirements are needed to facilitate additional IPv6 T&E.

Criterion 10: Demonstrate tactical deployability and ad hoc networking.

- Improvements in mobile applications have been demonstrated (autoconfiguration, multicasting, MANET protocols), but much work remains for development and T&E of the tactical deployability and ad hoc networking capabilities of IPv6.

- Mobility applications (NEMO and MANET) are in general an emerging technology. IPv6 T&E for this criterion are dependent on standards and mobile applications development.

4 Recommendations

Based on T&E results, analyses, and DoD Component's input, the following recommendations are made for full demonstration of the Joint Staff IPv6 operational criteria and to ensure a smooth transition to IPv6 for the DoD.

Although considerable T&E was accomplished for IPv6 security (Criterion 1) this reporting period, commercial development and implementation of security devices/applications is still needed to demonstrate this criterion. Recommendations include:

- Acquire pre-production HAIPE v3 devices, conduct beta T&E in both mixed IPv4/IPv6 and native IPv6 environments, and provide performance feedback to vendors.

- Perform vulnerability analysis, and formulate mitigation and configuration guidance for IPv6 implementations. Continue IPv6 T&E efforts for routers, switches, and security products.

- Emphasize the requirement for full IPv6 IPSec implementations, specifically in host OSs.

- Develop and conduct T&E of IPv6-capable Authentication, Authorization, and Accounting (AAA) and the PKI infrastructure.

- Collaborate with the National Information Assurance Partnership to develop protection profiles for the certification of IPv6 security products.

IPv6 interoperability (Criterion 2) and performance (Criterion 3) are expected to be fully demonstrated within the next Fiscal Year. Elements of transition mechanisms (Criterion 8) related to network transition (versus application transition) have already been successfully demonstrated. However, more experience using mixed IPv4/IPv6 networks in an operationally-realistic environment is needed. IPv6 should be deployed in selected Milestone Objective 2 (MO2) environments (as described in the MTP v2.0).

Voice, data, and video integration (Criterion 4) and operation in low-bandwidth environment (Criterion 5) both require policies, requirements, technical guidelines, and defined standards to demonstrate these criteria. Additional recommendations include:

- Encourage vendors to develop RTP and SIP (i.e., AS-SIP) products.

- Conduct T&E in operationally-realistic environments to demonstrate operations in low-bandwidth environments.

Mobility (Criterion 7), network management (Criterion 9), and ad hoc networking (Criterion 10) lacked development and implementation. Recommendations include:

- Develop a comprehensive T&E plan addressing elements to be demonstrated; this would serve as a guide to focus both DoD and vendor efforts.

- Direct vendors to use the DISR IPv6 Standard Profiles for IPv6 Capable Products for product development and implementations.

5 Summary

Products, applications, and standards critical to the DoD's IPv6 transition are still in development. Commercial availability of IPv6-capable security products (e.g., HAIPE v3 devices, firewall appliances, IDS, PKI functionality, and key distribution systems) that meet the DoD's IA requirements continues to be a major transition risk factor. Finally, IPv6 T&E and operational deployment of IPv6 capabilities may be delayed until the necessary standards, applications, and devices are commercially available.

Appendix A. References

- Public Law 109-163 National Defense Authorization Act for Fiscal Year 2006, January 6, 2006. http://www.defenselink.mil/dodgc/olc/docs/PL109-163.pdf

- Public Law 108-375 National Defense Authorization Act for Fiscal Year 2005, October 28, 2004. http://www.defenselink.mil/dodgc/olc/docs/PL108-375.pdf

- Department of Defense Internet Protocol Version 6 Master Test Plan, Version 2.0, September 2006. https://gesportal.dod.mil/sites/DoD_IPv6/IPv6_Documents/DoD_IPv6_TE_Report/2006.09.29_DoD_IPv6_MTP_V2.pdf

- Department of Defense (DoD) Internet Protocol Version 6 (IPv6) Transition Plan Version 2.0, June 2006. https://gesportal.dod.mil/sites/DoD_IPv6/IPv6_Documents/DoD_IPv6_Transition_Plan/2006.06.30_DoD_IPv6_Transition_Plan_V2.pdf

- DoD IPv6 Generic Test Plan, Version 2, September 2006. https://gesportal.dod.mil/sites/JITCIPv6/tewg/Document%20Library/1/IPv6%20Generic%20Test%20Plan/IPv6-GTPv2.pdf

Appendix B. Terms and Definitions

Approved Products List (APL): A registry of information technology products which have been assessed by DoD entities and have passed DoD interoperability and information assurance (IA) requirements.

Demonstration: Testing that is limited to a combination of related, perhaps interdependent, features or functions. It is usually an ordered sequence of tasks and is restricted from any operational network traffic.

DoD Components: The Office of the Secretary of Defense, Military Services, Chairman of the Joint Chiefs of Staff, Combatant Commands, Office of the Inspector General of the Department of Defense, Defense Agencies, DoD Field Activities, and all other organizational entities in the Department of Defense.

Engineering Analysis: Category of testing based on engineers' previous experience with the technology, as well as use of equipment specifications to speculate about the performance or capability.

Exercise: Testing that uses an operationally-realistic network with controlled traffic and realistic loading. Focus is on network and communications testing and includes automated test generators to assess the devices or systems functionality and performance. The DoD APL testing is also included in this category.

Experiment: Testing that consists of a scope that is restricted to a single question or theory with a test network isolated from operational network traffic. Few repetitions of test cases and a limited number of participants are involved.

Field Test: Testing that uses an operationally-realistic network with common protocol traffic and assumed loading conditions. Focus is on the devices or systems operating within the environment in which it is deployed. A well-defined, limited duration is set for testing.

IPv6 capable: An IPv6-capable system or product shall be capable (once IPv6 enabled) of receiving, processing, and forwarding IPv6 packets and/or interfacing with other systems and protocols in a manner similar to that of IPv4.

IPv6 Generic Test Plan Version 2 (IPv6 GTP): A plan developed to specify conformance, interoperability, and performance procedures that IPv6 products must successfully complete in order to be certified for interoperability by DISA (JITC).
http://jitc.fhu.disa.mil/adv_ip/register/register.html

Joint Staff IPv6 operational criteria: Criteria that must be successfully demonstrated to support a decision to initiate DoD transition to IPv6 and identify key operational and technical capabilities at a high level.

Milestone Objective 2 (MO2): DoD Components are authorized to implement and operate IPv6 across cooperative domain boundaries. At MO2, the policies, procedures, and technical guidance have been developed to expand the operation of IPv6 across cooperative domain boundaries, but limited to within DoD networks (no internet exchange of IPv6 packets, native or tunneled). MO2 will provide the ability to evaluate the scalability and further evaluate the IPv6 IA implications using tunneling and native IPv6 routing, as available. IPv6 traffic which crosses cooperative domain boundaries must be approved in accordance with the DISN connection-approval process to ensure compliance with IA policies. Multiple certification and accreditation authorities may be involved in MO2. MO2 permits applications to test IPv6-specific end-to-end capabilities and routing schema efficiencies. Limiting operation to within DoD, and only at approved locations, reduces risk to IA and operational impacts on existing IPv4 networks. MO2 was authorized as of October 1, 2006.

Mixed IPv4 and IPv6 environment: A mixed IPv4 and IPv6 environment includes the situations of tunneling IPv4 over IPv6 native network, tunneling IPv6 over an IPv4 native network, providing protocol translation at various points, and dual-stack operation.

Modeling and Simulation (M&S): Testing that uses a completely virtual environment to predict system or network performance. Software is used to simulate all involved devices and protocols.

Pilot: Testing that uses a functional, operational network with a limited number of administrators and users, but is realistic for the size of the network. There is no set time limit in conducting pilots and all traffic is non-scripted (routine traffic).

Appendix C. Acronym List

A	DNS A record for an IPv4 Address
AAA	Authorization, Authentication, and Accounting
AAAA	DNS AAAA record for an IPv6 Address
ACL	Access Control List
AFIOC	Air Force Information Operations Center
AFIWC	Air Force Information Warfare Center
AFNAS	Air Force Network Architecture Solutions
AFSN	Air Force System Networking
AH	Authentication Header
AIPTL	Advanced IP Technology Laboratory
ALG	Application Layer Gateway
AODV	Ad Hoc On-Demand Distance Vector
APL	Approved Products List
ARP	Address Resolution Protocol
AS-SIP	Assured Services-SIP
ASD	Assistant Secretary of Defense
ATM	Asynchronous Transfer Mode
BGP	Border Gateway Protocol
BIA	Bump in the Application Programming Interface
BIND	Berkeley Internet Name Domain
BIS	Bump In the Stack
BSD	Berkeley Software Distribution
CA	Certificate Authority
CDS	Cross Domain Solutions
CEF	Cisco Express Forwarding
CERDEC	Communications-Electronics Research, Development, and Engineering Center
CHS	Common Hardware System
CIO	Chief Information Officer
CLAN	Coalition LAN
CN	Correspondent Node
COI	Community of Interest
CONUS	Continental United States
COTS	Commercial Off The Shelf
CPU	Computer Processor Unit
CT	Cipher Text
DAD	Duplicate Address Detection
DHCP	Dynamic Host Configuration Protocol
DHCPv6	Dynamic Host Configuration Protocol version 6
DiffServ	Differentiated Services

DISA	Defense Information Systems Agency
DISN	Defense Information Systems Network
DISR	DoD IT Standards Registry
DITO	DoD IPv6 Transition Office
DNS	Domain Name System
DoD	Department of Defense
DoS	Denial of Service
DREN	Defense Research and Engineering Network
DS-3	Digital Signal Level 3
DSCP	DiffServ Code Point
DUT	Device Under Test
DVBS	Digital Video Broadcast-Satellite
EAP	Extensible Authentication Protocol
ESP	Encapsulating Security Payload
FIPS	Federal Information Processing Standard
FTP	File Transfer Protocol
FY	Fiscal Year
GB	GigaByte
GBS	Global Broadcast Service
GES	Ground Entry Sites
GIG	Global Information Grid
GN	Ground Node
GOTS	Government Off The Shelf
GPS	Global Positioning System
GRE	Generic Routing Encapsulation
GTP	Generic Test Plan
HA	Home Agent
HAIPE	High Assurance Internet Protocol Encryptor
HF	High Frequency
HTTP	Hypertext Transfer Protocol
HPCMP	High Performance Computing Modernization Program
I3MP	Installation Information Infrastructure Modernization Program
IA	Information Assurance
IATF	IA Technical Framework
ICE	IPv6 Capable Exercise
ICMP	Internet Control Message Protocol
ICMPv6	Internet Control Message Protocol version 6
ID	Identification
IDS	Intrusion Detection System
IE	Internet Explorer
IETF	Internet Engineering Task Force

IIS	Internet Information Services
IKE	Internet Key Exchange
IM-PEPD	Implicit Peer Enclave Prefix Discovery Protocol
INSC	Interoperable Networks for Secure Communications
IOS	Internetwork Operating System
IOZ	Future Capabilities Division
IP	Internet Protocol
IPS	Intrusion Prevention System
IPSec	IP Security
IPTV	Internet Protocol Television
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
ISATAP	Intra-Site Automatic Tunneling Address Protocol
IT	Information Technology
JCAN	Joint Capability for Airborne Networking
JGN	Joint Gateway Node
JITC	Joint Interoperability Test Command
JTEN	Joint Tactical Edge Networks
JUICE	Joint User Interoperability Communications Exercise
JVMF	Joint Variable Message Format
K	Kilobit
Kbps	Kilobits per second
L2	Layer 2
L3	Layer 3
L3G	Multicast L3 Gateway
LAN	Local Area Network
LDAP	Lightweight Directory Access Protocol
M&S	Modeling and Simulation
MAC	Media Access Control
MANET	Mobile Ad hoc Networks
Mb	Megabit
Mbps	Megabits per second
MIB	Management Information Base
MCS	Maneuver Control System
MIP	Mobile IP
MIPv6	Mobile IP version 6
MN	Mobile Node
MO2	Milestone Objective 2
MR	Mobile Router
MTPv2.0	Master Test Plan Version 2
NAT	Network Address Translation

NEMO	Network Mobility
NIAP	National Information Assurance Partnership
NII	Networks and Information Integration
NIPRNet	Unclassified-But Sensitive IP Router Network
NM	Network Management
NM/OPS	NM Operations
NMS	Network Management Systems
NOC-V	Network Operations Center – Vehicle
NS	Name Server
NS	Neighbor Solicitation
NSA	National Security Agency
NTP	Network Time Protocol
OC	Optical Carrier
OLSR	Optimized Link State Routing
OPNET	Optimized Network Evaluation Tool
OS(s)	Operating System(s)
OSPF	Open Shortest Path First
OSPFv3	Open Shortest Path First version 3
OTM	On The Move
PC	Personal Computer
PIC	Physical Interface Card
PKI	Public Key Infrastructure
PO	Participating Organization
POP3	Post Office Protocol version 3
PPP	Point-to-Point Protocol
PS	Policy Servers
PT	Plain Text
QFY	Quarter Fiscal Year
QoS	Quality of Service
RA	Router Advertisement
RDT&E	Research, Development, Test, and Evaluation
RDP	Remote Desktop Protocol
RF	Radio Frequency
RFC	Requests for Comment
RIM	Radio Interface Module
RIP	Routing Information Protocol
RO	Route Optimization
RSVP	Resource Reservation Protocol
RTP	Real Time Protocol
RTSP	Real Time Streaming Protocol
S&TCD	Space and Terrestrial Communications Directorate

S/A/C	Services/Agencies/Components
SATSIM	Satellite Simulator
SDP	Service Delivery Points
SDP	Shelf Discovery Protocol
SEND	Secure Neighbor Discovery
SIIT	Stateless IP/Internet Control Message Protocol Translation
SIP	Session Initiation Protocol
SMTP	Simple Mail Transfer Protocol
SNMP	Simple Network Management Protocol
SP	Service Pack
SPAWAR	Space and Naval Warfare Systems
SSHv2	Secure Shell Version 2
STIG	Secure Technical Implementation Guide
3DES	Triple Data Encryption Standard
T&E	Test and Evaluation
TCP	Transmission Control Protocol
TDM	Time Division Multiplexer
TEWG	Test and Evaluation Working Group
TGA	Traffic Generator/Analyzer
TIC	Technology Integration Center
TOC	Tactical Operation Center
TRT	Transport Relay Translator
UDP	User Datagram Protocol
UHF	Ultra High Frequency
URL	Uniform Resource Locator
USAISEC	U.S. Army Information Systems Engineering Command
v	Version
VLAN	Virtual Local Area Network
VLG	Virtual / Live Gateway
VoIP	Voice over IP
VPN	Virtual Private Network
WAN	Wide Area Network
WIN-T	Warfighter Information Network-Tactical
WLAN	Wireless LAN
WWW	World Wide Web

Appendix D. DoD IPv6 2007 Test Report Summaries

This appendix provides summaries for the 44 IPv6 T&E reports that DoD Components submitted for this year. The applicability of each report to the Joint Staff IPv6 operational criteria is summarized in Table D-1. The alphanumeric designator that precedes each report title in this table corresponds to the section number of the appendix that summarizes the report. Each report summary is comprised of the following eight elements: title, testing organization and publication date, summary, T&E method, relevant Joint Staff IPv6 operational criteria (including Level 1 and 2 decomposition relevancy), configuration, results, and conclusions/recommendations.

Table D-1 2007 T&E Reports and Related Operational Criteria

Section	Test Report Title	Joint Staff IPv6 Operational Criteria									
		1	2	3	4	5	6	7	8	9	10
D.1	DoD IPv6 Transition Office (DITO) IPv6 Domain Name System (DNS) Test Report		X						X		
D.2	Warfighter Information Network-Tactical Internet Protocol Version 6 Assessment		X			X			X		
D.3	Defense Research and Engineering Network IPv6 Introduction	X	X				X		X		
D.4	Defense Research and Engineering Network IPv6 Lessons Learned	X	X				X		X		
D.5	High Assurance Internet Protocol Encryptor (HAIPE) Program Overview Summary 0.0.1	X									
D.6	JUICE 2006 IPv6 Transition Mechanism Test Report V 2.0								X		
D.7	JUICE 2006 IPv6 Transition Mechanism Test Report V 2.0: Appendix A Draft 0.9, IPv6 Transition Mechanism Alternatives Study: Maneuver Control System Proof of Concept								X		
D.8	IPv6 Security Assessment	X									
D.9	Joint Users Interoperability Communications Exercise 2006 Internet Protocol Version 6 Information Assurance Assessment Report	X							X		
D.10	IPv6 MO1 Test Report for IPv6 Security Concerns	X							X		
D.11	Global Broadcast Service (GBS) Integration with IPv6, a Pilot Implementation			X					X		
D.12	Multi-Level Security, Geographically Targeted Information Dissemination Using Internet Protocol Version 6 (IPv6)	X									
D.13	INSC Task 3 (Mobility) Final Report	X						X	X		X
D.14	INSC Test and Demonstration Architecture for INSC Phase II	X	X		X				X		X
D.15	GIG-EF Event 06-3 IPv4 and IPv6 Security Hop-by-Hop Control Plane Tests	X							X		
D.16	Internet Protocol Version 6 Joint Staff Operational Criteria 2 and 3, Phase I Test Report		X	X					X		
D.17	Special Interoperability Test Certification of the Juniper M and T Series Routers for IPv6 Capability	X	X						X		
D.18	Defense Research and Engineering Network Juniper ISG-2000 Firewall Test Report	X							X		
D.19	Internet Protocol Version 6 Low Bandwidth Test Report		X	X		X					
D.20	IPv6 Autoconfiguration White Paper										X
D.21	Operational Issues with IPv6 DNS		X						X		
D.22	IPv6 Multihoming White Paper										X

Table D-1 2007 T&E Reports and Related Operational Criteria (continued)

Section	Test Report Title	Joint Staff IPv6 Operational Criteria									
		1	2	3	4	5	6	7	8	9	10
D.23	Special Interoperability Test Certification of Microsoft Windows Vista Enterprise Operating System installed on a Panasonic Toughbook CF-74 and a Panasonic Toughbook CF-51 for Internet Protocol Version 6 (IPv6) Capability	X	X						X		
D.24	Special Interoperability Test Certification of Techguard PoliWall for Internet Protocol Version 6 (IPv6) Capability	X							X		
D.25	IPv6 Protocol Security Assessment and Issues	X						X	X	X	X
D.26	JCS Criteria 4, Phase I – Demonstration of QoS Capabilities of IPv6 Using DiffServ (FY07 Moonv6 Demonstration)				X	X			X		
D.27	Test Results and Lessons Learned	X		X					X		
D.28	Network Management IPv6 Feasibility Study Report							X			
D.29	IPv6 Vulnerability Assessment Report for the Air Force Standard Desktop Configuration for Microsoft Windows Vista	X							X		
D.30	IPv6 Scalability Testing Final Report						X		X		
D.31	Milestone Objective 2 IPv6 Scenario 1 Implementation Guide & Test Parameters	X							X		
D.32	MO2 Security Concerns for Microsoft Windows IPv6 Protocol	X	X						X		
D.33	Milestone Objective 2 IPv6 Scenario 1 Router Configuration Guide	X							X		
D.34	2006 Ethernet Switch Comparison Report	X	X	X					X	X	
D.35	Implicit Peer Enclave Prefix Discovery Protocol (IM-PEPD) High Assurance Internet Protocol Encryptor (HAIPE) Discovery White Paper	X									
D.36	JUICE 2006 Test Report Verification of IPv6 Stateless Auto-configuration, Tactical Reorganization & Network Mobility (NEMO)										X
D.37	Implementing Internet Protocol Version 6 (IPv6) on an Army Installation	X	X						X		
D.38	Juniper Networks Internet Protocol Version 6 Report			X					X		
D.39	Beyond Addresses: IPv6 Value for the GIG	X	X		X						
D.40	Testing Known Vulnerabilities Against Internet Protocol Version 6 (IPv6)	X							X		
D.41	Internet Protocol Version 6 (IPv6) Mitigation Planning Phase 3: Custom Configuration Guidance	X							X		
D.42	Internet Protocol Version 6 (IPv6) Mitigation Planning Phase 4: RFCs and Protocols	X							X		

Table D.1 2007 T&E Reports and Related Operational Criteria (continued)

Section	Test Report Title	Joint Staff IPv6 Operational Criteria									
		1	2	3	4	5	6	7	8	9	10
D.43	Internet Protocol Version 6 (IPv6) Mitigation Planning Phase 1: Vendor Patch Implementation Plan	X							X		
D.44	Internet Protocol Version 6 (IPv6) Mitigation Plan Phase 2: STIG Implementation	X							X		
Total Test Reports by Joint Staff IPv6 Operational Criteria		**29**	**14**	**6**	**3**	**3**	**3**	**3**	**34**	**2**	**6**

D.1 DoD IPv6 Transition Office (DITO) IPv6 Domain Name System (DNS) Test Report

Testing Organization and Publication Date

DITO/SI International
11 July 2006

Summary

The DITO IPv6 DNS Test Report examined which versions of Microsoft Windows DNS could experience issues during the transition of the .mil domain to a dual-stack primary master name server model. This report focused on Iterative Mode Resolver functionality after the transition of the primary master name server to dual-stack operations.

Test and Evaluation Method

Experiment

Joint Staff Operational Criteria Tested

2 (2.1, 2.1.1, 2.2, 2.2.1)
8 (8.1, 8.1.1)

Configuration

This test bed consists of a single DNS server and active directory domain. The primary master name server for this domain is authoritative (as opposed to the Internet Service Provider or name registration company maintaining authority) for the domain. The primary domain controller (TRUMAN) for the active directory domain is also a DNS. The addition of a Windows 2000 and Windows NT4 server provides the breadth of testing desired. Table D-2 describes the servers and network equipment utilized in testing.

Table D-2 Equipment Configuration

Server Name/Equipment	Platform	Software
NS1	Fedora Core 4	BIND 9.3.1
Truman	Windows 2003 Server	DNS SP1
WIN2KDNS	Windows 2000 Server	DNS SP3
NT4DNS	Windows NT4 Server	SP6+
Cisco	2900 Switch	Not Listed
Cisco	3550 Switch	Not Listed
Cisco	2600 Router	Not Listed
Cisco	PIX 515e Firewall	Not Listed
Juniper	M71 Router	Not Listed

Results

After installation of the OS and installation of patches and services packs, the DNS service was started. The server was configured as a secondary server for the ipv6lab.com domain. The DNS service was then started. Test results are summarized below.

- Unknown Record Type Testing for NT4
  - Systems were still responsive. A set of queries to the server showed it was operational, responding to a DNS query for an IPv4 address (A) and a DNS query for an IPv6 address (AAAA) queries over IPv4 transport.

- Unknown Record Type Testing for Windows 2000
  - The system did not halt. A set of queries to the server showed it was operational, responding to A and AAAA queries over IPv4 transport.

- Unknown Record Type Testing for Windows 2003
  - The system did not halt. A set of queries to the server showed it was operational, responding to A and AAAA queries over IPv4 transport.

- Wrong Transport Testing for Windows NT4 Server
  - Microsoft does not provide a supported or experimental IPv6 stack for NT4. Therefore, testing of the wrong transport issue was unavailable on Windows NT4 Server.

- Wrong Transport Testing for Windows 2000 Server
  - An unsuccessful attempt was made to install the IPv6 stack on Windows 2000. Therefore, testing of the wrong transport issue was unavailable on Windows 2000 Server. Note that the IPv6 stack is experimental and not supported by Microsoft.

- Wrong Transport Issue for Windows Server 2003
  - Systems were still responsive. A set of queries to the server showed it was operational.

Conclusions

No tested version of Windows DNS service displayed the unknown record type issues. Windows 2003 DNS service did not display the wrong transport issues. Windows NT4 and Window 2000 DNS services were not tested for the wrong transport issue because supported IPv6 stack implementations were not available. With the use of Windows 2003 DNS, the potential issues related to a dual-stack root server and primary master name servers are eliminated.

D.2 Warfighter Information Network-Tactical Internet Protocol Version 6 Assessment

Testing Organization and Publication Date

JITC
August 2006

Summary

For the JUICE 2006, JITC assessed the WIN-T's core communicating protocols operating in IPv6. Protocols included H.323 VoIP, HTTP, FTP, SMTP, and POP3. The JITC's Advanced IP Technology Laboratory (AIPTL) assessed the WIN-T's Joint Gateway Node (JGN) at Fort Huachuca, Arizona, and Taunton, Massachusetts, from 19 July to 10 August 2006. Personnel from JITC, General Dynamics, and the WIN-T program office were involved in testing.

Test and Evaluation Method

Exercise

Joint Staff Operational Criteria Tested

2 (2.3, 2.3.1, 2.3.2)
5 (5.1, 5.1.1)
8 (8.1, 8.1.1)

Configuration

The test network consisted of two terrestrial links: a 512 Kbps traffic channel (test traffic only) and a 256 Kbps management channel. The management channel was built so that the Agilent blade at the AIPTL could control and manage the remote blade in Taunton, Massachusetts. The original management channel was programmed for 64 Kbps. However, due to unforeseen downloads initiated by the blades upon a reboot, that link was increased to 256 Kbps to speed up the download process. Table D-3 displays the tested equipment platform and associated software versions.

Table D-3 Equipment Configuration

Equipment (#)	Platform	Software
WIN-T Border Router	Cisco 3745	12.3
WIN-T LAN Router	Cisco 3745	12.3
JITC AIPTL Router	Cisco 3845	12.3
JITC AIPTL Router (2)	Juniper M40e	7.4R2.6
JITC AIPTL Router (2)	Juniper T320	7.4R2.6
JITC AIPTL Router	Cisco 3845	12.4
JITC AIPTL Switch	Cisco 6500	12.2
JITC NIT Lab Router	Cisco 3745	12.4

Results

3,263 of 3,265 calls were successfully completed, for a 99.93% success rate. These calls varied from 30 seconds to 8 hours, using the G.711 codec (the Defense Switched Network standard). All 700 data transfers were exchanged error-free. This included 340 HTTP, 120 FTP, 120 SMTP, and 120 POP3 transfers. Table D-4 lists the results of each scenario.

Table D-4 WIN-T Test Results

Scenario	Duration (seconds)	Voice		Data								Percent Complete
				HTTP		FTP		SMTP		POP3		
		Comp	Att	Comp	Att	Comp	Att	Comp	Att	Comp	Att	
150 Voice Only	30	150	150									100
200 Voice Only	30	200	200									100
250 Voice Only	30	250	250									100
300 Voice Only	30	300	300									100
350 Voice Only	30	350	350									100
400 Voice Only	30	400	400									100
5 Voice, Data	30	5	5	5	5	5	5	5	5	5	5	100
20 Voice, Data	60	20	20	10	10	10	10	10	10	10	10	100
40 Voice, Data	60	40	40	10	10	10	10	10	10	10	10	100
100 Voice, Data	3600	100	100	10	10	10	10	10	10	10	10	100
50 Voice, 50 HTTP	30	50	50	50	50							100
100 Voice	30	100	100									100
100 Voice, 100HTTP	30	100	100	100	100							100
100 Voice, 20 HTTP	30	100	100	20	20							100
100 Voice, 50 HTTP	30	100	100	50	50							100
100 Voice, 10 Data	60	100	100	10	10	10	10	10	10	10	10	100
100 Voice, 100 Data	28800	100	100	10	10	10	10	10	10	10	10	100
50 Voice, 10 Data	30	50	50	10	10	10	10	10	10	10	10	100
100 Voice, 10 Data	30	100	100	10	10	10	10	10	10	10	10	100

Table D-4 WIN-T Test Results (continued)

Scenario	Duration (seconds)	Voice		Data								Percent Complete
				HTTP		FTP		SMTP		POP3		
		Comp	Att	Comp	Att	Comp	Att	Comp	Att	Comp	Att	
150 Voice, 10 Data	30	100	100	10	10	10	10	10	10	10	10	100
150 Voice, 10 Data	120	148	150	10	10	10	10	10	10	10	10	98.9
200 Voice, 15 Data	30	200	200	15	15	15	15	15	15	15	15	100
200 Voice, 10 Data	30	200	200	10	10	10	10	10	10	10	10	100
Total		3263	3265	340	340	120	120	120	120	120	120	99.9

Conclusions

The JGN demonstrated the ability to support core communication protocols in IPv6. This includes H.323 VoIP, HTTP, FTP, SMTP, and POP3. It is recommended however, that the JGN and other WIN-T components undergo further IPv6 testing. As IPv6 progresses, other key areas of emphasis such as security, mobility, and ad hoc networking should be examined for compatibility within WIN-T.

D.3 Defense Research and Engineering Network IPv6 Introduction

Testing Organization and Publication Date

High Performance Computing Modernization Program (HPCMP)/OSD
July 2006

Summary

The DREN IPv6 pilot peers with the Internet, commercial, and other DoD networks in attempt to demonstrate security, interoperability, scalability, and transition mechanisms within a live network to support DoD transition to IPv6. This document introduces the architecture behind the IPv6 pilot network, including node, link, network, and security details.

Test and Evaluation Method

Engineering Analysis

Joint Staff Operational Criteria Tested

1 (1.1, 1.1.1, 1.5, 1.5.1)
2 (2.1, 2.1.1, 2.1.3, 2.2, 2.2.1, 2.2.3, 2.3, 2.3.1, 2.3.3)
6 (6.1, 6.1.4)
8 (8.1, 8.1.1)

Configuration

The DREN configuration consists of ten core nodes on Optical Carrier (OC)-192c backbone Continental United States (CONUS), with OC-12c extensions to Alaska and Hawaii. Approximately 120 sites [Service Delivery Points (SDP)] are also connected from Digital Signal Level 3 (DS-3) to OC-48c. Numerous devices and software support DREN services, but it is a predominately-unclassified network with some Type 1 encryptors.

Results

One of the best reasons why the DREN IPv6 pilot has been developed is that it has given the DoD community a production environment to more directly test a functional network, as opposed to a closed, limited test network. It is also assisting other DoD agencies in configuration, management, security, and deployment of an IPv6 network. The lessons learned and research to be conducted for this effort will greatly benefit the DoD community.

Some of the current accomplishments of the DREN IPv6 Pilot are an enabled IPv6 Wide Area Network (WAN) infrastructure, security and performance equivalency to IPv4, facilitation of IPv6 deployment to HPCMP funded sites' infrastructure, equipment, and lessons learned feedback.

Conclusions

The DREN IPv6 pilot has provided great research to the DoD community concerning areas of security, interoperability, scalability, and transition mechanisms within a live IPv6 network. This effort is providing the community an insight on possible problems, advice on implementation order, and guidance in selecting a transition technique that is feasible and effective.

D.4 Defense Research and Engineering Network IPv6 Lessons Learned

Testing Organization and Publication Date

HPCMP/OSD
July 2006

Summary

The DREN IPv6 pilot peers with the Internet, commercial, and other DoD networks in an attempt to demonstrate security, interoperability, scalability, and transition mechanisms within a live network to support DoD transition to IPv6. The lessons learned from this document give other S/A/C guidance on IPv6 implementation and helps them avoid mistakes.

Test and Evaluation Method

Pilot

Joint Staff Operational Criteria Tested

1 (1.1, 1.1.1, 1.5, 1.5.1)
2 (2.1, 2.1.1, 2.1.3, 2.2, 2.2.1, 2.2.3, 2.3, 2.3.1, 2.3.3)
6 (6.1, 6.1.4)
8 (8.1, 8.1.1)

Configuration

The DREN configuration consists of 10 core nodes on OC-192c backbone (CONUS), with OC-12c extensions to Alaska and Hawaii. Approximately 120 sites SDPs are also connected at DS-3 to OC-48c. Numerous devices and software support DREN services, but it is a predominately unclassified network with some Type 1 encryptors.

Results/Lessons Learned

There are four primary goals of the DREN IPv6 pilot: IPv6 enabled WAN infrastructure, security and performance, facilitate IPv6 deployment into HPCMP funded sites' infrastructures, and IPv6 enabled. The lessons learned in attaining each goal will assist the DoD and other S/A/C to transition to IPv6.

- Goal 1: IPv6 enabled WAN infrastructure
  - Need more extended use in "real" IPv6 networks to expose and fix remaining errors
  - Query vendors for specific features that matter; their interpretation may differ
  - Memory may become an issue for dual-stack support
  - Many routers have fairly complete production-quality IPv6

- Most products (90%+) claim IPSec support for IPv6, but are not functionally complete
- A large percentage of routers built since 2001 handle IPv6 fairly well
- Static routing is not scalable; hard to maintain more than 12 sites
- Long-term routing protocol solutions: internal Border Gateway Protocol (BGP), external BGP and OSPF Version 3 (OSPFv3)
- Transition mechanisms should be used sparingly (current limitations)
- DREN has experience in using Tunnel Broker (commercial Hexago works well) and open source 6to4 tunnel software (works with limitations inside one enclave but not recommended across enclaves)

- Goal 2: Security and Performance
  - Internally developed an update to IDS software used to support IPv6 as well as IPv4; no commercial sources for wide-area IDS available. This updated IDS is available to other S/A/C within DoD
  - Firewalls are slowly becoming available; current devices have limitations
  - Lack of scanning tools
  - Standards and implementations of IPSec remains a problem
  - Security was a priority for the DREN IPv6 pilot, but documenting it has been difficult

- Goal 3: Facilitate IPv6 deployment into HPCMP funded sites' infrastructures
  - Procurement
    - Most commercial computers are hardware capable
    - Recommend upgrade/replace router >4 years old
  - People
    - Surprisingly little training required for technical personnel
    - Attitude adjustment needed for security, procurement, and management
  - Process
    - Incorporate IPv6 support considerations in system support, planning, and installation processes
    - Software upgrades and network/system reconfiguration
  - The order in which networks, computers, applications, and DNS are IPv6 enabled impact transition. Recommend:
    - Networks, DNS software, and other IP infrastructure
    - Computers: servers first and desktop later
    - Applications: clients first and server software later
    - Make DNS entry changes last (no right time to change DNS).

- Goal 4: IPv6 enabled
  - At each site, IPv6 transition was done by a small number of part-time technical personnel, as an additional duty, and without any additional funding
  - This caused a lack of detailed tracking of expenses and summary data
  - Site trends:
    - Few purchases were necessary
    - Common to expand memory on routers

        - No computer replacement
        - Upgrades on OSs at no cost under standard maintenance contract with exception of Microsoft OS
        - Training: commercial, HPCMP provided, and self

- Additional Lessons Learned
    - Network Management has the fewest software tools with IPv6 capability compared to all functional areas
    - A template to follow for technology transition is much better than starting from scratch.

Conclusions

The DREN IPv6 pilot provided great research to the DoD community concerning areas of security, interoperability, scalability, and transition mechanisms within a live IPv6 network. Many issues remain within the IPv6 transition; however, great strides have been made in developing an operational IPv6 network. The lessons learned provide the community an insight on possible problems, advice on implementation order, and guidance in selecting a transition technique that is feasible and effective.

D.5 High Assurance Internet Protocol Encryptor (HAIPE) Program Overview Summary 0.0.1

Testing Organization and Publication Date

NSA
August 2006

Summary

A HAIPE is a programmable IP Information Security device with traffic protection, networking, and management features that provide IA services for IPv4 and IPv6 networks. HAIPEs that are v3 compliant meet the DoD mandate for IPv6 compatibility and the goals of the Cryptographic Modernization Initiative, and are a key component of the Global Information Grid (GIG) Vision. This document explains the core requirements of HAIPEv3 and compares current features to previous versions.

Test and Evaluation Method

Engineering Analysis

Joint Staff Operational Criteria Tested

1 (1.6, 1.6.1)

Configuration

The HAIPEv3 is expected to be backward compatible to HAIPE Interoperability Specification (IS) 1.3.5. It is also planned to be the baseline for network encryptors deployed within the GIG.

Results

HAIPEv3 is expected to offer many benefits over the current HAIPE IS 1.3.5. It is expected that HAIPEv3 will provide services that allow communities to meet their traffic protection, networking, and management needs. Four goals of HAIPEv3 include:

- Bandwidth Efficiency
  - Reduce encapsulation overhead
  - Reduce cryptographic transform overhead.

- Over-the-Network Management
  - Management Information Base Version 3.0 (MIBv3.0)
  - SNMP Version 3
  - Firmware download.

- Signaling Interoperability
  - HAIPE to HAIPE
  - HAIPE to infrastructure
  - HAIPE to Key Management Infrastructure
  - HAIPE to Response Service Message

- HAIPE Implementations
  - Enclave Gateway
  - Host
  - Terminal

Conclusions

This document provides insight on the development or/and acquisition of HAIPEv3. As the primary network encryptor for the deployment of future GIG networks, HAIPEv3 must provide adequate traffic protection, networking, and management features.

D.6 JUICE 2006 IPv6 Transition Mechanism Test Report V 2.0

Testing Organization and Publication Date

Space and Terrestrial Communications Directorate (S&TCD), Communications-Electronics Research, Development and Engineering Center (CERDEC), Software Engineering Center (SEC), Communications-Electronics Life-Cycle Management Center (CE-LCMC), Software Engineering Directorate (SED)
October 2006

Summary

As part of a multi-phase study of IPv6 transition mechanisms, the Army funded the development of an IPv4-to-IPv6 Transformer by Datatek Applications, Inc., prototyped an Application Layer Gateway (ALG), Transport Relay Translator (TRT), and Multicast L3 Gateway (L3G) based on legacy Maneuver Control System (MCS), and tested a commercially-available IPv4-to-IPv6 Tunnel Broker by Hexago. This report examined the functionality and interoperability of five IPv6 transition mechanisms in a tactical, operational, satellite communications network during the JUICE 2006 event.

Test and Evaluation Method

Exercise

Joint Staff Operational Criteria Tested

8 (8.1, 8.1.1)

Configuration

This test consisted of nine scenarios demonstrating several modes of communications between various MCS configurations. The configurations include IPv4-only, IPv6-only, dual stacked with the ALG and TRT, IPv4-only using the Datatek Transformer, and IPv6-only using the Hexago Tunnel Broker.

This test utilized routers, computers, satellite communications, and additional network equipment. Table D-5 lists the equipment and descriptions.

Table D-5 Equipment Configuration

Equipment	Description/Software
Router 1	Compaq PC with FreeBSD, dual stack with FreeBSD, dual stack with IPv6 multicast routing
Router 2	Compaq PC with FreeBSD, dual stack with L3G
Computer	Dell PC with Windows XP and MCS, IPv4-only
Computer	Dell PC with Windows XP and MCS, both IPv6 only
Computer	Dell PC with Windows XP and MCS, both dual stack
Computer	Dell Laptop with Windows XP and MCS, both IPv4-only
Computer	Dell Laptop with Windows XP (dual stack) and MCS (IPv6-only)
Datatek Transformer	N/A
Hexago Tunnel Broker	N/A
Netgear Ethernet Switches	N/A
Memotec CX960e	Satellite Router Gateway for IPv4 link
L3 MDL8372S0004	Satellite Transmit Modulator for IPv4 link
L3 MDL8471F0008	Satellite Receive Demodulator for IPv4 link
Comtech EF Data CDM570L	Satellite Modem for IPv6 link

Results

Application Level Gateway

The ALG functions seamlessly and requires no additional administrative effort beyond the existing baseline. Neither legacy software function nor user operation is adversely affected by the ALG. Although MCS provides two unicast modes, UDP and TCP, the ALG addresses only TCP unicast.

Transport Relay Translator

The TRT operates automatically and works well in conjunction with the Datatek Transformer and Hexago Tunnel Broker since it makes no differentiation between message sources. An MCS node's ability to access the TRT is not limited by any additional transition mechanism. Enabling the TRT on a dual-stack MCS node requires only the configuration of a single environment variable and no additional administrative effort.

Datatek Transformer

The Datatek Transformer provides unicast connectivity between a single IPv4 node and an IPv6 network. It must be configured separately, but does not require client software or additional configuration of the IPv4 node. It lacks support of some network features, such as multicast, but Datatek is continuing to improve the product.

Hexago Tunnel Broker

The Hexago Tunnel Broker can provide either IPv4 or IPv6 tunneling. In this exercise, the Hexago Tunnel Broker client was connected to an IPv4-only network and required IPv6-in-IPv4 tunneling. On start up, the client established a tunnel to the Tunnel Broker server over the IPv4 satellite link. This tunnel provided IPv6 network access to the client as if it were on a native IPv6 network. To function properly, the tunnel broker server and client (the "endpoints" of the tunnel) must be dual stacked. Usually, the client software is used on a host node, but it may also be implemented on a network router, providing access to the tunnel to several users at once.

Multicast Gateway

The multicast gateway must be implemented on a network router running Protocol Independent Multicast for IPv6, Sparse Mode (PIM6-S), for which Cisco has implemented support only very recently in Version 12.4 of its OS. While it is the preferred solution, it was not available in time for this exercise. The multicast gateway used in this study was an experimental version that runs only on Free Berkeley Software Distribution (BSD). The multicast gateway requires no additional administration of network nodes. It is not limited to MCS, but is compatible with all multicast messaging.

Conclusions

The transition mechanisms tested for this exercise performed the functions for which they were designed in both laboratory and tactical satellite environments. While no single system can provide everything, when used correctly, individually or in concert, these systems provide efficient, inexpensive, and reliable connectivity in a complex networking environment. The MCS upgrades, the Datatek Transformer, and the Hexago Tunnel Broker all provided reliable transition support in specific circumstances.

D.7 JUICE 2006 IPv6 Transition Mechanism Test Report V 2.0: Appendix A Draft 0.9, IPv6 Transition Mechanism Alternatives Study: Maneuver Control System Proof of Concept

Testing Organization and Publication Date

S&CTD/CERDEC
March 2006

Summary

The CERDEC, S&TCD, and Telcordia were tasked to demonstrate the transition of a legacy software application and verify its operation within a hybrid network comprised of a legacy IPv4 system and a future network IPv6 system. The MCS was chosen as the software application, and the Network Operations Center - Vehicle (NOC-V) was chosen as the legacy system. An Optimized Network Evaluation Tool (OPNET) modeling and simulation (M&S) environment was used to represent a future force network. S&TCD provided system engineering and integration for the task and implemented the future force model. Telcordia performed a software analysis of the MCS and developed an ALG to enable the MCS to function in both IPv4 and IPv6 modes.

Test and Evaluation Method

M&S

Joint Staff Operational Criteria Tested

8 (8.1, 8.1.1)

Configuration

Using OPNET version 11.0, S&TCD created an IPv6 M&S environment that was used to represent an Army future force backbone network such as WIN-T. For this demonstration, a unique concept of using a Virtual / Live Gateway (VLG) to interface the legacy network with the M&S environment was introduced. The model for the demonstration consisted of two intermediate router nodes and two VLGs to allow bidirectional traffic to flow in and out of the model.

The VLG is essentially an Ethernet card that resides on the computer that is running the OPNET model. It is the physical interface between the live environment and the virtual environment. However, no such product exists, not even by OPNET. Therefore, S&TCD created one by modifying an Ethernet driver applet. The driver accepts TCP and UDP IPv6 packets and then inspects the packet for proper IPv6 header format and content. If satisfactory, the driver then triggers a token or message sequence in OPNET.

Results

The task culminated with a proof of concept exit demonstration consisting of six scenarios:

- Sustain IPv4 Legacy Baseline Interoperability
- Send and Receive IPv6 Joint Variable Message Format (JVMF) Messages
- Exchange IPv6 / IPv4 JVMF Message via TRT
- Multi-Destination Unicast JVMF Messages in a 6/4 Hybrid Environment
- Multicast JVMF Messages in a 6/4 Hybrid Environment
- v6-over-v4 Automatic Tunnel Broker.

All scenarios ran flawlessly and the demonstration was successful in conveying the message that the IPv6 transition of legacy components, at least in the case of the NOC-V and MCS, is relatively straightforward.

Conclusions

The VLG used for integrating the live network with the simulated OPNET future force network was an innovative idea that shows much promise for modeling IPv6 networks of any size. It can be a useful tool for analyzing and testing conformance, performance, and interoperability.

The results of this task were very encouraging for those concerned about the transition and impact of IPv6 on legacy applications and systems. While MCS represents only one program among hundreds in use by the Army, the experience with MCS enforces the general belief that applications written with modularity in mind, and that follow the concept of a layered model as advocated by the Open Systems Interconnection (OSI) reference model, are relatively easy to modify for forward compatibility with IPv6. In addition, it was found that a typical Army tactical communication system, such as the NOC-V, is capable of processing the throughput of IPv6 traffic by upgrading or replacing its L3 routing components. However, the Layer 2 (L2) routing components, i.e., Ethernet switches, are not impacted by the presence of IPv6. This is especially good news since most network topologies today use many more switches than routers. This applies to Army tactical networks as well, where two routers, one red and one black, may reside in a front-end communication system, such as the NOC-V or a Brigade Subscriber Node.

D.8 IPv6 Security Assessment

Testing Organization and Publication Date

Science Applications International Corporation (SAIC)
01 June 2006

Summary

This paper surveys the current availability and maturity of IPv6 across a range of products, from OS to utilities to security tools, and details many attacks on various parts of the IPv6 protocol suite. We assume non-IPSec use of the protocol, as will be common for some time. IPSec, while subject to its own attacks and defenses, usually mitigates the protocol attacks mentioned here. Solving the keying problem for large full mesh IPSec networks is left as an exercise for the reader. Organizations already implementing a PKI should be able to leverage it to ease implementation of IPSec. Testing was performed from May 2005 to October 2005.

Test and Evaluation Method

Experiment

Joint Staff Operational Criteria Tested

1 (1.1, 1.1.1,1.4, 1.4.1, 1.4.2)

Configuration

The following client OSs were tested: Fedora Linux Core 3 and 4, Solaris 9 x86, FreeBSD 5.4, OpenBSD 3.7, Windows XP Professional, Windows Server 2003, Longhorn 5213 and Longhorn 5270. Among these OSs, there were four IPv6 stack variants in use.

The Scapy4 packet construction toolkit was used to perform all protocol testing and attacks. At the start of this project, IPv6 support in Scapy4 was rudimentary. Over the course of this project, several contributors have improved Scapy4's IPv6 support.

Results

All client OSs tested properly rejected Neighbor Solicitation (NS) packets with a hop count less than 255. This eliminates threats of NS being used for discovery or Denial of Service (DoS) from off the local link of the target.

Linux, FreeBSD, Solaris 9, OpenBSD 3.6, and Windows Vista (Longhorn) were immune to attempts at denying IPv6 nodes their self-assigned addresses by sending out malicious Neighbor Advertisements in response to Duplicate Address Detection (DAD) packets at boot time. Linux, FreeBSD, and OpenBSD all logged a duplicate address error to syslog and continued to use the address. This was true for both the link local address and the global unicast addresses.

However, this attack was successful for Windows XP Professional and Windows Server 2003. Further, Windows XP and 2003 were susceptible to this attack at any time, not just during the window between when the DAD NS packet is sent and the DAD process times out waiting for a response. This means that any Windows XP or 2003 machine can be denied all IPv6 service by an attacker on the same link at any time.

Vendor OSs showed varying behavior in response to Router Advertisement attacks. For volume-based DoS attacks (e.g., receipt of a steady stream of 65536 unique router advertisements), Windows Server 2003 and Windows XP Professional both consumed all available CPU resources while processing the router advertisement packets. (Note that when the same sets of router advertisements were repeated, excessive CPU resources were not consumed.)

Fedora Core 3/4 is partially immune to the gratuitous Router Advertisements (65 Kilobit total) attack. First, when receiving the advertisements, far fewer CPU resources were consumed. Second, Fedora has a default upper limit of 16 addresses per interface that it will process. IPv6 stack implementers may want to add a "limit per second" for the number of Router Advertisements (RAs) that will be accepted.

FreeBSD and OpenBSD appear to lay between Windows and Fedora in their susceptibility to the gratuitous RA (65 K) attack. While the BSDs lacked an upper limit on the number of RAs accepted, they consumed fewer CPU resources than Windows, but more than Fedora when processing the requests.

Conclusions

While limited to on-link attacks, malicious RAs can deny service and provide for man-in-the-middle eavesdropping and packet injection. Allowing control of node interfaces and routing table entries via unauthenticated Internet Control Message Protocol (ICMP) Version 6 (ICMPv6) packets puts users and networks at risk. It is recommended that users and network administrators avoid the use of RAs in favor of Dynamic Host Configuration Protocol(DHCP) Version 6 (DHCPv6) or static configuration. While DHCPv6 itself is vulnerable to similar attacks, these attacks are more limited in scope and require timing. In addition, DHCP has been around longer and is better understood by administrators. Automatic configuration by DHCPv6 will also spare users and administrators from having to manually configure lengthy, complex IPv6 addresses.

D.9 Joint Users Interoperability Communications Exercise 2006 Internet Protocol Version 6 Information Assurance Assessment Report

Testing Organization and Publication Date

JITC
November 2006

Summary

The JITC performed an IA vulnerability assessment during JUICE 06. The JUICE 06 IA Assessment identified IPv6 IA vulnerabilities in individual devices and within networks that are representative of operational DoD systems. The JUICE 06 network architecture consisted of equipment and networks that were assessed within a simulated Defense Information Systems Network (DISN) Core. In addition to the IA assessment, the MO2 enclave was assessed based on DITO IA Guidance for MO2.

Test and Evaluation Method

Exercise

Joint Staff Operational Criteria Tested

1 (1.1, 1.1.1, 1.4, 1.4.1, 1.4.2, 1.5, 1.5.1)
8 (8.1, 8.1.1, 8.1.3)

Configuration

Testing consisted of equipment and networks that were assessed within a simulated DISN Core using one Juniper T640 and two Juniper T320s. The assessment also included the simulated Unclassified-But-Sensitive IP Router Network (NIPRNet) using one Juniper T640, two Juniper T320s, two Juniper M40es, and two Cisco 3845s and Teleport environments, which included NIPRNet plus the satellite simulator and one Cisco 3845. Both NIPRNet and Teleport environments used client resources provided by Microsoft Vista Beta Builds 5520 and 5472 and servers employing the Microsoft Longhorn Server Build 5520. The test network included three Windows XP clients, two Windows Beta clients, a Root Server, three Domain Controller Servers, and two member servers. The separate MO2 enclave assessment used two Windows XP clients, two Cisco 3745s, one Cisco 3725, and a NetScreen 500 firewall.

Results

The threat rating for the network assessed was calculated by dividing the sum of discovered findings by the total of all vulnerabilities checked. Table D-6 provides the numeric key to Table D-7. Table D-7 provides the threat rating averages for all of the assessed network components (i.e., routers and switches).

Table D-6 Combined Test Team IA Vulnerability Threat Rating Scheme

Threat Rating	Definition
High = 2 - 3	Highest total average of vulnerabilities.
Medium = 1 - 1.99	Medium total average of vulnerabilities.
Low = 0 - 0.99	Lowest total average of vulnerabilities.

The higher the threat rating in table D-7 given for each IA Technical Framework (IATF) area and subset, the greater the number of vulnerabilities discovered. When future assessments are performed on these networks, these results may be used for comparison. A comparison of implementation trends (by IATF area) will enable testers to see its posture improvement from assessment to assessment.

Table D-7 Combined Test Team IA Threat Ratings of Mission Critical Components

IATF Area	Network Component	Threat Rating
Network and Infrastructure	Routers	2.2
	Switches	1.7
Local Computer Environment	Windows Servers	1.5
	Windows Workstations	0.7

In the MO2 enclave, IPv6 traffic was transmitted from the IPv6 only (Cisco 3745) end of the MO2 enclave to the IPv4 only (Cisco 3725) end of the MO2 enclave. An Ethereal packet sniffer was set up in the enclave between the I1.B and the firewall. No IPv6 packets were detected on the IPv4 only side of the dual-stack router.

Although there was no IPv6 traffic found passing through the router, there was one significant finding. If a piece of equipment that has an IPv6 configuration is introduced on the IPv4 only network side without removing the IPv6 configuration, the piece of equipment installed would introduce IPv6 packets on the IPv4 network.

For the Build 5472, the test was between two Vista workstations in different domains. The tests were only run in a dual IPv4/IPv6 environment. In this Vista build, tunneling worked and a secure connection was setup and maintained during the assessment.

For the Build 5520 using the same setup as in Build 5472, there was communication between the workstations, but IPv6 IPSec tunneling was unable to be set up and the connection initiates and remains unsecured.

Conclusions

The vulnerabilities found in the IPv6 protocol are the same vulnerabilities known to exist in the IPv4 protocol. The SAINT, Nmap, and ThreatEx test tools are able to detect known IPv6 vulnerabilities.

Microsoft Vista Build 5472 supports IPSec tunneling and secure connection in a dual IPv4/IPv6 environment. Microsoft Build 5520 supports nonsecure communication, but does not support IPSec tunneling.

The MO2 enclave did not pass IPv6 traffic to the IPv4 side of the network. However, all IPv6 configurations must be removed on equipment inserted on the IPv4 side prior to sending packets across the network.

D.10 IPv6 MO1 Test Report for IPv6 Security Concerns

Testing Organization and Publication Date

Future Capabilities Division (IOZ) Air Force Information Warfare Center (AFIWC)
22 May 2006

Summary

The tests explained in this report were derived from the DISA DITO IA Interim Guidance for MO1 document. The results are recommend configurations for Air Force boundary protection and internal control devices to protect against IPv6 attacks. The IOZ Assessments Branch IPv6 team researched how IPv6 could affect an operational Air Force network, in order to provide confidence that IPv6 enabled nodes will not compromise security of the operational network. The test results presented in this report provide important information on the behavior of IPv6 in a predominately IPv4 environment.

Test and Evaluation Method

Experiment

Joint Staff Operational Criteria Tested

1 (1.4, 1.4.2, 1.5, 1.5.1)
8 (8.1, 8.1.1, 8.1.3)

Configuration

The following table shows the configuration of the IPv6 Lab's routers, computer systems, and test equipment used. The initial configuration of the routers involved in testing were determined by using the Air Force Ports, Protocols, and Services Matrix document release 3.0, February 28, 2005.

Table D-8 Equipment Configuration

Equipment (#)	Platform	IOS/Software
Router	Cisco 7206VXR	12.4(4)T1
	Cisco 2621	12.2(13)T16
Computer Systems	(2) Dell Dimension 4500	MS Windows XP Pro/Service Pack 2
	Dell Dimension 4300	MS Windows XP Pro/Service Pack 2
Server	Dell 350 Server Blade	MS Windows 2003 Enterprise/ Service Pack 1
Test Equipment	IXIA 400T	LM1000T5
	Ethereal	0.10.13(C)

Table D-8 Equipment Configuration (continued)

Equipment (#)	Platform	IOS/Software
Test Equipment	NetWag	5.33
	Kiwi Syslog Daemon	7.2.35

The Air Force Network Architecture Solutions (AFNAS) Lab was used as a third party test element to validate selected test objectives. Their equipment configuration is listed below.

Table D-9 AFNAS Equipment Configuration

Equipment (#)	Platform	IOS/Software
Routers	Cisco 7204	12.3-12a
	Cisco 7505	12.3-12a
Computer Systems	(2) Dell Optiplex GX620	MS Windows XP Pro/Service Pack 2
	Dell Latitude 505	MS Windows XP Pro/Service Pack 2
Server	Dell Power Edge 1750	MS Windows 2003 Enterprise/ Service Pack 1
Test Equipment	Ethereal	0.10.13(C)
	NetWag	5.33

Results

Filtering

The bulk of the test required that IPv6 traffic and specific IPv4 ports and protocols be denied at the IPv6 enclave boundary. The majority of these filter requirements mandated only a simple Cisco router Access Control List (ACL). The following table identifies each test objective, the method tested, and the result for each objective.

Table D-10 Test Results

Filter Objective	Test Method	Method	Result
Deny native IPv6 packets	Packet injection	ACL	Denied
Deny IPv6 in IPv4 tunnel/IPv4 protocol 41 packets	Packet injection	ACL	Denied
Deny IPv4 in IPv4 tunnel/IPv4 protocol 4 packets	Packet injection	ACL	Denied
Deny Source Demand Routing Protocol/IPv4 protocol 42 packets	Packet injection	ACL	Denied
Deny AX.25 tunnel/IPv4 protocol 93 packets	Packet injection	ACL	Denied
Deny IP-within-IP Encapsulation Protocol/IPv4 protocol 94 packets	Packet injection	ACL	Denied
Deny EtherIP/IPv4 protocol 97 packets	Packet injection	ACL	Denied
Deny Encapsulation Header /IPv4 protocol 98 packets	Packet injection	ACL	Denied
Deny Layer 2 Tunneling Protocol/IPv4 protocol 115 and UDP 1701 packets	Packet injection	ACL	Denied
Deny Generic Routing Encapsulation/IPv4 protocol 47 packets	Packet injection	ACL	Denied

Table D-10 Test Results (continued)

Filter Objective	Test Method	Method	Result
Deny Fragmented Packets	Packet injection	ACL	Denied
Deny IP Security/IPv4 Protocols 50 and 51 and UDP port 500	Packet injection	ACL	Denied
Teredo Transition Mechanism/IPv4 UDP port 3544	Packet injection	ACL	Denied

Administrative Information

There is a possibility of IPv6 administrative and control information traversing the enclave perimeter using IPv4 services. This information needs to be prevented from adversely affecting the enclave. Such traffic includes Internet Control Message Protocol (ICMP), DNS, SNMP, and routing protocol exchanges and updates. To minimize the effect, utilizing Air Force Ports, Protocols, and Services Matrix (PPS) filters to restrict access to management information to authorized systems and users where possible is recommended. In particular, the IPv6 Helper Service on Windows XP and 2003 family servers should be disabled as it exploits the propagation of DNSv6 addresses across an IPv4 protocol layer or 6in4 encapsulation.

AFNAS Test Results

The following table presents the five MO1 objectives selected for third party testing through the utilization of the AFNAS Lab.

Table D-11 Test Results

Filter Objective	Test Method	Method	Result
Teredo Transition Mechanism/IPv4 UDP port 3544	Packet injection	ACL	Denied
Deny Fragmented packets	Packet injection	ACL	Denied
Deny IPv6 in IPv4 tunnel/IPv4 protocol 41 packets	Packet injection	ACL	Denied
IPv6 Administrative Information	Disable specified service	IPv6 Helper Service	Eliminated AAAA record propagation
Deny native IPv6 packets	Packet injection	ACL	Denied

Testing within the AFNAS Lab verified the initial testing of IPv6 met the required MO1 objectives.

- Cisco Router ACL filters are capable of providing the necessary elimination of identified ports and protocols

- Modification of identified services will eliminate leakage of IPv6 administrative information at the Windows client and server system level.

Conclusions

The tests performed in the AFNAS and the IPv6 Labs proved that modifications/additions to existing ports and protocol filters are sufficient to satisfy MO1 IPv6 enclave requirements. Every specific identifiable port and protocol was denied at the enclave boundary router using a Cisco ACL; the inclusion of an additional firewall at the IPv6 enclave border is unwarranted. With the additional filters in place, the enclave devices were able to communicate with enterprise network devices over authorized ports and protocols.

The IDS approved for use by the Air Force does not offer full support of IPv6 alerts and/or monitoring. Until an extensive IPv6-supportable proxy and IDS/Intrusion Prevention System (IPS) are designated by the Air Force, full integration of IPv6 into the Air Force's Enterprise network will not be accomplished. The effort now is to control the IPv6 traffic propagated within a specific enclave.

Future testing will isolate any IPv6 traffic propagated across the enterprise layer with some type of tunnel where the traffic can either be encrypted using IPSec or controlled using ACLs.

D.11 Global Broadcast Service (GBS) Integration with IPv6, a Pilot Implementation

Testing Organization and Publication Date

DISA
2006

Summary

The GBS is a satellite communications program modeled on the highly successful commercial Digital Video Broadcast-Satellite (DVB-S) platform. It provides a worldwide, high-bandwidth, one-way transmission of classified and unclassified video, imagery, and other files to support joint military forces. Recently, GBS architecture has transitioned from Asynchronous Transfer Mode (ATM) technology to IP to enhance a multitude of features, such as modularity and decreased operational complexity. In support of the Office of Assistant Secretary of Defense (OASD)-NII mandate and the DoD's Network-centric vision, the GBS Joint Program Office and DISA have been jointly investigating the transition of GBS to IPv6. This report will disclose the results of performance metrics crucial to video and file dissemination.

Test and Evaluation Method

Demonstration

Joint Staff Operational Criteria Tested

3 (3.1, 3.1.1, 3.2, 3.2.1, 3.3, 3.3.1)
8 (8.1, 8.1.1, 8.1.2)

Configuration

The simulated operational testing architecture is composed of a DVB-S transport equipment string that encapsulated and modulated incoming IP packets onto a Radio Frequency (RF) carrier. To accurately simulate a satellite communications system, testing included a Satellite Link Emulator and Carrier-to-Noise Generator as well as an ingress and egress router. Spirent and Ixia gear were used for traffic loading and measuring performance.

Results

The GBS IPv6 pilot simulated operational testing to measure the performance and capability of the GBS DVB-S system to transport native IPv6 and dual-stack traffic. The performance of delivering IPv4 unicast traffic over the pilot test bed served as a comparison baseline. The Spirent SmartBits traffic generator was used to generate IPv4 and IPv6 unicast traffic at various frame sizes and data rates (test data rates: 2, 3, and 6 Megabits per second (Mbps) with frame sizes of 512, 1024, 1280, and 1400 Bytes for each data rate). The represented dual-stack traffic was composed of 95% IPv4 traffic and 5% IPv6 traffic, which is the projected ratio of IPv4 to IPv6 traffic during initial deployment of IPv6 over the GBS networks.

Frame Loss

The frame loss data demonstrates minimal loss for IPv4 unicast traffic, native IPv6 unicast traffic, and dual-stack traffic. The frame loss for all three traffic types was below .0005%, which is well within the acceptable limit of data loss across a satellite system. CPU processing on the network devices within the GBS IPv6 pilot architecture was minimal, indicating that operating in either native IPv6 or dual-stack mode will not affect the performance of GBS operational networks.

Latency

The lowest average latency was observed over the GBS pilot test bed for IPv6 unicast traffic, while the highest average latency was observed for IPv4 unicast traffic. For dual-stack traffic, the IPv6 packets arrive significantly in advance of the IPv4 packets. This illustrates the efficiency of the routers and network devices within the pilot network to process and forward IPv6 traffic. This is attributed to the streamlined design of the IPv6 header as compared to IPv4.

Conclusions

The performance of native IPv6 traffic and dual-stack traffic over the GBS IPv6 pilot architecture has proven to be more efficient than the current IPv4 architecture. Based on this initial data, deploying IPv6 and dual stack within the GBS architecture will not affect the performance of the network. In fact, deploying IPv6 within the GBS network will enhance the delivery of video and files by reducing latency across the network. Additional IPv4, IPv6, and dual-stack multicast testing will be completed over the GBS pilot simulated operational test bed to demonstrate the benefits of deploying IPv6 multicasting.

D.12 Multi-Level Security, Geographically Targeted Information Dissemination Using Internet Protocol Version 6 (IPv6)

Testing Organization and Publication Date

SI International and RGII
2006

Summary

This report specifically describes a security architecture that employs IPSec enhancements, flow labels and QoS, along with Global Position System (GPS) protected by level 1 encryption, to provide targeted information dissemination over the emerging DoD IPv6 network infrastructure.

Test and Evaluation Method

Engineering Analysis

Joint Staff Operational Criteria Tested

1 (1.1, 1.1.1)

Configuration

The IPv6 protocol contains many enhancements of IPv4 that can strengthen security of hosts and networks. Such features include:

- AH
- ESP
- Flow Labels.

Results

IPSec

The security architecture for IP comprises the use of AH, ESP, and Internet Key Exchange (IKE). The IPv6 base protocol specification requires that all implementations of IPv6 must include extension headers to support IPSec AH and ESP. This requirement, along with security transmission of keys using IKE Version 2, provides an end-to-end secure channel for communication.

Both AH and ESP can be activated in either tunnel mode or transport mode. Tunnel mode provides security mechanisms for the IP protocol layer. Transport mode provides security mechanisms for the protocol layers above the IP protocol layer.

Cross Domain Solutions (CDS)

Current CDS architectures do not address many existing and future Communities of Interest (COI) requirements. Additionally there is no easy way to identify risk, prioritize mitigation activations, and too much dependence on high assurance sentinel systems deployed throughout the network.

A solution for CDS is to have a centralized enclave or enclaves. An IPv6 IPSec based Virtual Private Network (VPN) architecture may allow end-to-end transport mode IPSec between hosts inside and outside the enclave. The security association in this scenario is between two hosts and the traffic within the enclave(s) is cipher text. As a result, some provisions must be made to audit the traffic using such end-to-end encryption schemes.

In the new IPv6 IPSec based VPN architecture, the solution management model would include a centralized administrative domain with an enclave policy and protection center allowing for more efficient Certification and Accreditation procedures.

Policy Servers

Central client Policy Servers (PS) can be utilized within the IPv6 CDS network architecture to authenticate network users as part of the log-on process. The centrally stored client policies can then be downloaded to the end system, such as PKI certificates.

Deployment Architecture

The multi-level security, geographically-targeted information dissemination architecture relies on an end-to-end security model employing a distributed PS authentication mechanism and IPv6 transport to achieve mobile, secure and authenticated CDS. The use of IPv6 flow labels adds additional granularity to the distribution of multicast information flows. These technologies can facilitate the timely dissemination of targeted tactical and situation information to a highly mobile warfighter while meeting information security and access authentication requirements.

Conclusions

This multi-level security, geographically-targeted information disseminations architecture is based on existing international standards based technologies. This architecture satisfies current DoD requirements concerning information security, access authentication, CDS and mobility. Further, it facilitates the efficient and timely distribution of information from Combined Intelligence Center, theater and battle unit sources using a single certificate authority to achieve ubiquitous security.

D.13 Interoperable Networks for Secure Communications Task 3 (Mobility) Final Report

Testing Organization and Publication Date

INSC
14 July 2006

Summary

The INSC project is an international collaborative research and development activity between Canada, France, Germany, Italy, Norway, Netherlands, the United Kingdom, and the United States. The project's goals are to specify, implement, test, and demonstrate common technical architecture for interoperable secure networks with mobility extensions, using commercial technologies, products and solutions wherever possible. This report gives a brief overview of the various mobility technologies.

Test and Evaluation Method

Demonstration

Joint Staff Operational Criteria Tested

1 (1.3, 1.3.1)
7 (7.1, 7.1.2)
8 (8.1, 8.1.1)
10 (10.1, 10.1.1, 10.1.2, 10.2, 10.2.1)

Configuration

The test bed configuration varied for each test and country. Most participating organizations continued to use 802.11-based Wireless Local Area Network (WLAN) technology as a proof-of-concept wireless ad hoc link in simulated, emulated, and live test configurations. Those tests included MANET, MIPv6, and NEMO features. Table D-12 lists most of the equipment used in the INSC testing framework.

Table D-12 Equipment Configuration

Equipment (#)	Platform	IOS/Software
Routers	Cisco 3640	Experimental version 12.4
	Cisco 7513	Experimental version 12.3
	Cisco 2600	Not listed
Switch	Cisco 6506	6.1(1a)
Computer Systems	Linux Mandrake	10.1
	Linux	Kernel 2.6.11-1
Wireless Access Point	Orinoco AP-100	Not listed

Results

MANET

The MANET unicast protocols examined in some detail by T3 include Ad Hoc On-Demand Distance Vector routing (AODV), Optimized Link State Routing (OLSR) and some functional variants, and MANET extensions to OSPF (MANET-OSPF). In addition to performance testing with automated tools, various real applications were functionally tested over these MANET protocols. VoIP, video streaming, web browsing, and chat applications have been run over these networks successfully.

A direct comparison of the Boeing MANET-OSPF code and US Naval Research Lab OLSR was conducted on the US-developed Mobile Network Emulator operating over 2 Mbps 802.11b WLAN, with the purpose of evaluating the relative performance of each protocol. The scenario chosen was a 10-node network with a random motion model (identical for each protocol) and a traffic scenario of Many-to-One, with each node sending Multi Generator data to a gateway node.

The aggregate performance of MANET-OSPF and OLSR were very similar in this small network scenario. The rate plots (shown in original document) have only slight differences and the latency of both protocols is dominated by retransmission buffers in the WLAN cards. It is expected, and has been shown in simulation, that the performance difference of these protocols is more visible when the networks are scaled up to a larger number of nodes.

MIPv6

During INSC Phase II, Italy developed an alternative solution to the MIPv6 bootstrapping problem that, compared to those being devised by the IETF, has the advantage of enabling better control of mobility service and removing the need for a full Extensible Authentication Protocol (EAP) between the MN and Home Agent (HA) for MIPv6 authorization.

The central element of the architecture is the AAA server on the home domain, which interacts with both the MN and the selected HA to perform service authorization and configuration. The solution is applicable to any access network relying on EAP for user authentication and works with all EAP methods supporting the exchange of general-purpose information elements, in any form. Exploiting this capability, the MN and home AAA server can piggyback MIPv6 negotiation messages within the same EAP conversation used to carry out user authentication.

The proposed architecture allows the home domain to maintain a centralized management (on the AAA server) of the user profiles and the AAA procedures for any type of service, including Mobile IPv6. Moreover, the solution has the following advantages:

- Improves the reliability and performance of the Mobile IPv6 protocol, in that the HA to be dynamically assigned to the MN can be freely chosen among those that are closest to the user’s point of attachment, thus optimizing network usage and reducing the transfer delay for data traffic in bi-directional tunneling

- Can be deployed or extended with new features, without having to update the access equipment and the AAA protocols in use. Only minor changes in the AAA servers, the HAs and the mobile terminals are required; the AAA client does not play any active role in MIPv6 negotiation (i.e., it is a pass-through for EAP signaling). This reduces the deployment costs and makes the solution easy to use even when a MN is roaming with an administrative domain different from its own

- Allows the usage of any AAA protocol supporting the transport of EAP messages for the communication between the AAA client and server. This significantly simplifies the deployment of MIPv6 in existing communication networks, where support for Diameter protocol in access equipment is not so extensive

- Allows the home domain to dynamically choose the authentication method for IKE bootstrapping and to automatically distribute the pre-shared key eventually needed. In this way, the pre-shared key need not be preconfigured and can be frequently changed, increasing resistance to attacks. In the case of an EAP method providing dynamic generation of keying material, the pre-shared key can be derived from EAP hierarchy, avoiding the need to explicitly send it to the MN.

NEMO

Preliminary tests revealed a major issue with the Cisco 7500 platform: enabling the HA functionality jeopardized the forwarding capability of the router. Packets, including those not related to NEMO, were forwarded through the correct interface, but used an incorrect next hop Media Access Control (MAC) address. As a workaround, this platform was relegated to the role of a Mobile Router (MR).

MR handovers were emulated by implementing each access subnet as a separate Virtual LAN (VLAN) on a programmable switch and changing the VLAN associated with the MR egress interface port.

The goal of this functional test was to analyze the correctness and stability of the NEMO protocol when a single mobile network dynamically changes its point of attachment to the WAN. For this test, the MR1 continuously roamed among its home network and both foreign networks, while the Correspondent Node (CN) continuously sent ICMPv6 echo request packets to the MR1's home address.

By monitoring echo request and echo reply exchanges on the home network, it was verified that the NEMO Implementations (NEPL) correctly provided the expected mobility support. Additionally, the sniffing station was able to display the respective NEMO control information, such as Binding Updates and Binding Acknowledgements exchanged between MR1 and HA.

Two issues were found during NEMO testing: missing override flags in HA's proxy neighbor advertisement and an empty home agent list in Dynamic Home Agent Address Discovery.

Table D-13 presents the handoff latency measurements of NEMO testing.

Table D-13 Handoff Latency Measurements

Scenario	Average Handoff Latency (sec)	Standard Deviation (sec)	Median Handoff Latency (sec)
Cisco-Cisco	8.543	0.005	8.542
Nested Cisco-Cisco	10.544	0.005	10.545
Linux-Linux	2.392	0.166	2.380
Nested Linux-Linux	2.606	0.292	2.584
Linux MR - Cisco HA with DHAAD	4.261	0.183	4.272
Nested Linux MR - Cisco HA with DHAAD	4.524	0.355	4.413
Linux MR – Cisco HA without DHAAD	4.218	0.204	4.142
Nested Linux MR - Cisco HA without DHAAD	4.353	0.335	4.270

Conclusions

MANET technology is useful for supporting military network regions requiring self-organization, mesh operation, and possible high mobility. While different MANET protocols have different performance behaviors, and while a subset of two or three may cover most military scenarios, there is no "one size fits all" design at present. The appropriateness of different solutions has been shown to be related to the intended operational scenarios, applications, and platform requirements (e.g., proactive vs. reactive, vehicular vs. manpack, convoy vs. cluster at deployed Headquarters, local vs. non-local communications).

MIPv6 implementations have become more mature and more widely available, but the inclusion of IPSec integration into the protocols is still evolving. Standard specifications for Hierarchical Mobile IP (MIP) have advanced more slowly than anticipated at the beginning of INSC Phase II and have been studied less during this time period.

NEMO provides newer technology enhancement for aggregate prefix mobility. As a newer technology extension of MIPv6 concepts, NEMO may provide more relevant military support for larger platform and network mobility across and within WAN architectures. Task 3 investigated this technology's basic modes and some extended operational modes (e.g., NEMO nesting). Basic NEMO is presently functional, but implementations are still not stable (more testing is recommended). Software bugs and missing protocol features have been discovered by Task 3 researchers. The NEMO nesting function (multiple NEMO tunnels) has been shown to work, but there is little experience with the performance impact of current solutions. There are also important architectural issues resulting from additional encapsulation with each nesting level.

D.14 INSC Test and Demonstration Architecture for INSC Phase II

Testing Organization and Publication Date

INSC
13 December 2005

Summary

This report describes the test and demonstration architecture that will be employed during phase II of the INSC project. The proposed architecture has been developed to allow and facilitate test and demonstration of specific research topics specified within the four INSC II technical tasks: security, mobility, network and traffic management, and WAN and IPv4/IPv6 internetworking.

Test and Evaluation Method

Engineering Analysis

Joint Staff Operational Criteria Tested

1 (1.1, 1.1.1)
2 (2.1, 2.1.1, 2.1.2, 2.2, 2.2.1, 2.2.2, 2.3, 2.3.1, 2.3.2)
4 (4.1, 4.1.1)
8 (8.1, 8.1.1, 8.2, 8.2.1)
10 (10.1, 10.1.1, 10.1.2, 10.2, 10.2.1)

Configuration

Participants in this demonstration included Canada, France, Germany, Italy, the Netherlands, Norway, the United Kingdom, the United States, and the North Atlantic Treaty Organization Consultation, Command, and Control Agency. Each participating organization (PO) has its own unique architecture; however, all POs consisted of IPv6/IPv4 gateways, National or Coalition LANs, WANs, and IPSec for encryption.

Results/Analysis

Security

The focus of security within INSC II is to optimize the security elements within tactical networks. This optimization is done in the following directions and for the related reasons:

- Automatic discovery of active network security devices: a proposed new protocol (IPSec Discovery Protocol) allows a timely discovery of actual active IPSec devices within a coalition network; their ongoing supervision whether these devices are still active over the time; and the reporting of valid routing prefixes behind the IPSec devices towards the other connected red routing domains

- Secure multicast communications between connected red routing domains: the IPSec protocol specification is enhanced to support a complete multicast communication between several end users

- A dynamic group key generation and distribution system is installed (both for the Multicast Discovery Protocol and for the multicast IPSec protocol) to support changing group composition (both for the operation based entering and leaving of various groups and the active exclusion of multicast group members based on security considerations)

- Concept of a tactical PKI: tactical PKI is based on the concept that a couple of certification authorities (one per participating nation) are installed on the battlefield, all configured as a sub-ordinate Certificate Authority (CA) to the respective national (strategic) CA. The trust-relationship is achieved by exchanging and accepting the related root certificates between all participating nations prior to the operation of the coalition network

- Application based secure user communication: a prototype of a Secure Communications Interoperability Protocol terminal is provided, allowing a secure communication between users located in red network domains (single end users or users within command and control entities) and users outside a tactical network. For this purpose an upper layer protocol stack (initially for packetized voice communication) is provided (based on SIP). This allows alternatively an encrypted or unencrypted communication over IP and provides a security gateway between unencrypted black networks and IPSec protected red networks.

Mobility

For test and demonstration in INSC II, a variety of MANET routing protocols will be installed and tested on emulated or actual wireless routing nodes within the architecture. IP MANET routing protocol variants based on open specification work ongoing within the IETF (e.g., OLSR, AODV, MANET-OSPF) over an 802.11 or other identified radio(s) interfaces are planned. The MANET gateway routers will at least have one wireless interface for ad hoc routing and one wired interface (e.g., Ethernet) for external INSC connections. The planned use of 802.11 WLAN technologies to support MANET and MIPv6 operations by many of the participants is targeted as a “proof-of-concept” capability to test and demonstrate IPv6 and IPv4 MANET-enhanced mobile routing and user roaming capabilities. The networking solutions investigated are adaptable to multiple radio technologies and some POs involved in INSC II have plans to investigate additional wireless technologies within the architecture.

Network and Traffic Management

It is anticipated that future national and coalition operations will be conducted using communications services provided by multiple WANs which may be provided by any combination of coalition private networks, national private networks, tunnels through secure national networks and commercial networks. All coalition traffic crossing these networks will be secured by the use of IPSec gateways at the Coalition LAN (CLAN)/WAN boundaries.

The INSC II demonstrations will show how audio/video calls and video streaming sessions can be set up with guaranteed QoS and dynamic admission control. The proposed framework for QoS Control uses IETF concepts for policy-based management and bandwidth brokerage. QoS is provided in the WAN and CLANs, with at least three DiffServ traffic classes.

WAN and IPv4/IPv6 Networking

The INSC II Architecture contains two IPv6 WAN networks. Several nations will provide links between these two WAN networks allowing investigation into how traffic is routed between them. Additionally, investigation into how fragmentation or loss of connectivity between elements within the autonomous system affects how effectively traffic can flow end to end will be carried out.

Testing of this functionality will include the monitoring of multiple flows of traffic between the two WAN networks, looking at the path over which the traffic flows. Changes to the internal connectivity of one of the WAN networks will then allow testing of how BGP handles network fragmentation.

The INSC II Test and Demonstration Architecture consists of both IPv6 (WAN 1 and 2) and IPv4 (v4WAN) only networks, as well as dual-stacked networks (CLANs). The architecture employs (within the CLANs) both dual-stacked and single-stacked applications. It also includes an IPV4 to IPv6 gateway function. Using this architecture, it is hoped that all, or nearly all, of the transition scenarios identified above will be demonstrated. To investigate fully the implications of IPv4–IPv6 transition, the techniques that will be tested and demonstrated in INSC II will be employed not only individually but in some scenarios, “in tandem” (e.g., Translation-to-Tunneling-to-dual stack).

Conclusions

The INSC II focus on testing/analyzing security, mobility, network and traffic management, and WAN and IPv4/IPv6 internetworking in a coalition environment will be a commendable effort for testing and evaluating key components of the IPv6 transition.

D.15 GIG-EF Event 06-3 IPv4 and IPv6 Security Hop-by-Hop Control Plane Tests

Testing Organization and Publication Date

Space and Naval Warfare Systems (SPAWAR) Center San Diego
30 September 2006

Summary

This report details the control plane tests executed on the RoQ1 Linux router based architecture at the GIG-Evaluation Facilities. The tests examined the configuration overhead associated with the implementation of hop-by-hop control plane security in an IPv4 and IPv6 environment and the security implications of default configurations.

Test and Evaluation Method

Experiment

Joint Staff Operational Criteria Tested

1 (1.1, 1.1.1)
8 (8.1, 8.1.1)

Configuration

Test configuration included two Dell Dimension desktop computers configured with a Gnu's Not Unix/Linux OS with Fedora Core 5 kernel release 2.6.17-1.2141. A single link Linux router network setup with two routers using the Linux default authentication and encryption algorithms – Triple Data Encryption Standard (3DES) and Secure Hash Algorithm Version 1 (SHA1) were used. Wireshark (formally Ethereal) was utilized to monitor message passing on the connected link.

Results

Wireshark showed that, in contrast to Routing Information Protocol (RIP) and BGP routing protocols, the OSPF protocol does not use TCP or UDP, but rather IP datagrams directly. Hence, any non-unicast OSPF elements were not encrypted in IPv4/v6 default transport mode. The multicast OSPF Hello packets were the result of the routine Link State Database updates that were sent to both OSPF routers' addresses.

Another instance of an unencrypted message exchange on the local link occurs with the Address Resolution Protocol (ARP). The underlying fact that ARP resides below the IP layer allows all ARP traffic to remain unencrypted even with IPv4/IPv6 security enabled on the link. ARP requests were periodically directed upon the expiration of the ARP cache lifetime. Process resets also initiate the ARP request-reply exchanges.

Conclusions

The implementation of transport mode hop-by-hop control plane security only secures unicast messages on the link. Multicast messages, which are an integral component of the IPv6 standard, remain unencrypted. The multicast exchanges may be secured via the invocation of carefully configured access lists. The inherent complexity of some of the IPv6 mechanisms was made evident by the traces, where a significant amount of control plane traffic was transmitted unencrypted during initialization sequences as well as periodically. Further analysis of the threats associated with such an environment may be necessary to evaluate the actual risk versus the potential reward of absolutely securing the control plane.

D.16 Internet Protocol Version 6 Joint Staff Operational Criteria 2 and 3, Phase I Test Report

Testing Organization and Publication Date

JITC
May 2006

Summary

The 2006 IPv6 Joint Staff Operational Criteria 2 and 3, Phase I Test aimed to demonstrate end-to-end multi-protocol interoperability across a test network intended to simulate the GIG and to test IPv6 performance on individual network devices. The JITC Fort Huachuca, JITC Indian Head, SPAWAR San Diego, SPAWAR Charleston, and Air Force Communications Agency all participated in the exercise, which was conducted 6 through 22 November 2006.

Test and Evaluation Method

Exercise

Joint Staff Operational Criteria Tested

2 (2.1, 2.1.1, 2.1.2, 2.2, 2.2.1, 2.2.2, 2.3, 2.3.1, 2.3.2)
3 (3.1, 3.1.1, 3.2, 3.2.1, 3.3, 3.3.1)
8 (8.1, 8.1.1, 8.1.2)

Configuration

Interoperability

Tests were established to route multi-protocol traffic from multiple testing partners through the JITC Fort Huachuca GIG test node via the JITC Indian Head, GIG optical core. Results were collected at testing endpoints strategically placed to measure end-to-end multi-protocol traffic passing through the GIG and optical core.

Automated performance testing tools were used to generate concurrent, multi-protocol traffic (data, voice, and video) and report results for IPv4 and IPv6. Ixia Chariot (IxChariot) was used for end-to-end interoperability and performance testing. Scripts were subsequently combined into a single "Triple Play" test, consisting of data, voice, and video endpoint pairs. Identical "Triple Play" tests were developed for IPv4 and IPv6 to demonstrate multi-protocol interoperability.

Testing partners hosted one or more endpoints on various platforms with Ixia Chariot Endpoint Version 6.30 installed. Although traffic traversed the DREN, traffic flow was controlled via Generic Routing Encapsulation (GRE) tunnels, which bypassed normal DREN routing and allowed traffic control through the test network. Traffic passed between Fort Huachuca and

Indian Head via ATM Permanent Virtual Circuits, which provided the least obtrusive transport through the optical core and enhanced path realism.

Performance

Performance tests were executed in a closed network environment using automated test tools on GIG network devices, which in this case were GIG Core and Edge devices. Performance tests were run using the Spirent SmartFlow software on the SMB-600B SmartBits chassis to evaluate the compliance of a system or component with the performance requirements. The SmartFlow automated tool set tests the performance of an individual device under test (DUT) in a closed, controlled environment from a bit-loading standpoint. Identical tests were executed, first using IPv4, then IPv6. The router was configured with the initial frame size set to 64 and 76 bytes respectively for IPv4 and IPv6, then progressed through the following frame sizes (in bytes): 128, 256, 512, 1024, 1280, and 1518.

Results

Interoperability

Results for protocol-based interoperability tests are split into two categories dependent upon script function: transaction-based and streaming-based. Transaction based scripts resulted in 100% successful transport and delivery for all protocols and traffic types tested. The following scripts were transaction based:

- DNS
- FTP (Get/Put)
- HTTP (Text/Graphics Interchange Format)
- HTTP Secure (HTTPS)
- POP3
- SMTP
- SNMP

Streaming-based scripts successfully completed when running concurrently with transaction-based scripts with no significant loss. Aggregate streaming results for all sites tested resulted in 0.84% packet loss for IPv6 when compared to IPv4. The impact of packet loss was insignificant to overall function. The following scripts were streaming-based:

- G.711u (VoIP)
- IPTV - Video
- IPTV - Audio

Performance

All routers tested were configured for IPv4 and subsequently IPv6 processing with the frame sizes set to 64 and 76 bytes for small packets through 1518 bytes for large packets respectively.

Cisco 3745 Router - IOS 12.4(8a)

The device generated higher throughput on larger packets for both IPv4 and IPv6. The IPv6 latency was equivalent to IPv4 and throughput was equal to, or better than IPv4.

Cisco 3845 Router - IOS 12.4(1b)

Throughput was lower and latency higher on larger packet sizes, indicating an anomaly. No other symptoms were observed in the collected data and the router performed on par with IPv4 using different IOS versions.

Cisco 3845 Router - IOS 12.3(14)T2

The device generated higher throughput on larger packets for both IPv4 and IPv6. Latency increased for both IPv4 and IPv6 as the packet size increased. The IPv6 latency was generally higher than IPv4. However, IPv6 throughput was equal to, or better than IPv4.

Juniper M5 Router – OS 7.3R2.10

The traffic passed 100 Megabits (Mb) throughput for all IPv4 and IPv6 packets. The IPv6 latency was equivalent to IPv4 and throughput was equal to, or better than IPv4. Table 10 shows test results.

Juniper M40e Router - JUNOS 7.4R2.6

The traffic passed 100 Mb throughput for all IPv4 and IPv6 packets. The IPv6 latency was equivalent to IPv4 and throughput was equal to, or better than IPv4.

Conclusions

The interoperability test resulted in 99.48% successful transport, delivery, and interoperability for all transaction based protocols and traffic types tested in a mixed IPv4 and IPv6 environment. The test on streaming packets resulted in 99.16% of all protocols and traffic packets passing in the same mixed environment. The performance tests showed IPv4/IPv6 parity on the devices tested.

D.17 Special Interoperability Test Certification of the Juniper M and T Series Routers for IPv6 Capability

Testing Organization and Publication Date

JITC
February 2007

Summary

This report displays the results of the Special Interoperability Test Certification of the Juniper M and T Series Routers configured to support dual-stack IPv4/IPv6 protocols and the Adaptive Services Physical Interface Cards (PIC) necessary to support IPSec. Testing occurred from 11 September to 31 October 2006 at the AIPTL, JITC. Upon successful testing, the DUTs were placed on the DoD APL.

Test and Evaluation Method

Field Test

Joint Staff Operational Criteria Tested

1 (1.1, 1.1.1, 1.2, 1.2.1, 1.4, 1.4.1)
2 (2.1, 2.1.1, 2.2, 2.2.1, 2.3, 2.3.1)
8 (8.1, 8.1.1)

Configuration

The DUTs were part of a simulated DISN core test architecture managed by the AIPTL at JITC. The core consisted of one Juniper T640, two Juniper T320 routers, two Juniper M40e routers, two Cisco 3845 routers, and four Gateway workstations. One client was loaded with Windows XP, Silver Creek Pro SNMP Test Suite, Simple Tester Pro SNMP Test Suite, Ixia IxChariot Performance Tester, and an IP Packet capturing tool, Wireshark. The second client was loaded with Windows 2003 and Wireshark. The third workstation was loaded with Windows XP and was used to manage the Spirent SmartBits and ThreatEx test equipment. The fourth workstation was loaded with Windows XP and managed the Ixia Ix400T Traffic Generator/Analyzer (TGA).

Results

Core IPv6 Functionality

The hosts were able to send, receive, and process ICMPv6 packets and the routers processed the multicast requests by sending Echo Replies. A datagram of 1518 Kbps was transmitted from the TGA to the router. The router returned the correct “Packet Too Big” message and then refused to fragment the packet, subsequently dropping the packet. Analyzed network traffic also revealed a solicitation for a router subnet prefix, the client receiving the prefix, and sending a neighbor discovery. DAD was performed and no other device on the network possessed the address, and the client assigned the address to its main network interface. Application Traffic was sent through the network, which included TCP packets consisting of HTTP, FTP, and SMTP to the Juniper routers. The second test sent UDP packets consisting of Real-time Streaming Protocol (RTSP). The Juniper routers were able to establish, maintain, and terminate TCP and UDP connections across the network.

Routing and Switching Protocols

Multi-protocol BGP Extensions for IPv6, external BGP and internal BGP were configured on the M40e routers as per the current GIG architecture. The downstream Cisco Customer Edge device/transition router was configured with Multi-protocol BGP Extensions for IPv6, and configured as an External BGP Peer device. It was determined the DUTs were able to process the advertised routes and choose the correct path for the incoming packets.

Transition Mechanisms

All devices were configured with IPv6 and IPv4 TCP/IP stacks. The Spirent SmartBits was used to generate HTTP, FTP, SMTP, and RTSP traffic across the network. The DUTs were able to process all of the traffic types for both IPv6 and IPv4 TCP/IP stacks. Testers then manually configured IPv6 over IPv4 tunnels on two Juniper routers and one Cisco router. The Spirent SmartBits was used to simultaneously transmit IPv6 data packets (HTTP, FTP, SMTP, and RTSP) across an IPv4 tunnel. The Juniper routers were capable of processing the data packets through the IPv4 tunnels.

IA

For testing IPSec, IKE, and Internet Security Association and Key Management Protocol (ISAKMP), the Juniper routers were additionally configured with Adaptive Service PICs. The Spirent SmartBits generated IKE and ISAKMP traffic across a tunnel set up between the two M40e routers. The Juniper routers successfully sent secure traffic through the core network using the tunnel.

QoS

The proposed DoD use of DiffServ Code Points within the GIG dictated the setup parameters for this test. The DiffServ Code Point values were set on the two Cisco boundary routers of the simulated DISN test network. Bidirectional traffic was routed through the simulated DISN test network and out each of the boundary routers to their respective destinations. All DiffServ Code Points were correctly processed by each of the respective Cisco boundary routers. Juniper to Juniper DiffServ was tested successfully and Juniper was found to be compliant with RFC 2474.

Network Operations and Management

SNMP software testing tools (Simple Tester Pro v.10.0.3 and Silver Creek Pro v.10.4.7) were used to test RFCs 3411, 3412, 3413, and 4022. Both the Silver Creek Pro SNMP Vulnerability and Simple Tester Pro Performance Test Suites were used in addition to the RFC testing. The Juniper routers passed all required tests along with the Silver Creek vulnerability tests and the Simple Tester Pro performance tests.

Conclusions

The Juniper M and T Series Routers meet the IPv6-capable requirements and are certified for listing on the DoD APL as IPv6 capable. The DUTs successfully completed the related IPv6 performance and interoperability portions of the DoD IPv6 GTP Version 2.

D.18 Defense Research and Engineering Network Juniper ISG-2000 Firewall Test Report

Testing Organization and Publication Date

HPCMP
21 February 2007

Summary

The DREN IPv6 pilot conducted a test of the IPv6 capabilities of the Juniper Networks Corporation NetScreen ISG-2000 firewall. From March to October 2006, 10 locations connected to the DREN tested an unreleased (beta) version of the ScreenOS router OS on a production network.

Test and Evaluation Method

Pilot

Joint Staff Operational Criteria Tested

1 (1.4, 1.4.2, 1.5, 1.5.1)
8 (8.1, 8.1.1, 8.1.3)

Configuration

Testing consisted of 10 locations connected to the DREN operating in a dual-stack network infrastructure. The ISG-2000 firewalls were run both in-line and inserted in parallel with existing firewalls at the perimeter of the sites' infrastructure. IPv6 traffic was then routed through the ISG-2000 firewall while IPv4 traffic continued to be routed through the existing firewall. The ISG-2000 firewall was also intermittently operated in normal production mode on production networks during the test period.

Results//Lessons Learned

The ISG-2000 firewalls and beta ScreenOS software under test provide good but not perfect IPv6 screening capabilities.

Testing showed that IPv6 worked well on both the host side and the router side interfaces. The firewall was able to receive and deploy address space from the upstream router as well as hand out addresses to the hosts downstream.

Performance was on par with IPv4. Testing demonstrated that the firewall was able to adequately handle simultaneous stateful inspection of IPv4 and IPv6 data streams with little or no CPU impact. Performance measurements of the Saturn interface were taken under load. No buffer overruns or any other significant interface errors were recorded. As with their IPv4 code, automatic negotiations on their interfaces do not handle well with non-Juniper products.

Conclusions

Juniper ISG-2000 firewalls are suitable for use in IPv6 and IPv4/IPv6 (dual-stack) configurations. Testing to verify suitability for specific environments is still recommended (as is the case for IPv4 configurations).

D.19 Internet Protocol Version 6 Low Bandwidth Test Report

Testing Organization and Publication Date

JITC
June 2007

Summary

The JITC, Fort Huachuca, Arizona, conducted interoperability and performance testing from 27 November through 29 January 2007. The test event compared the interoperability and performance characteristics of IPv6 compared to IPv4 within a low-bandwidth Time Division Multiplexer (TDM) network environment.

Test and Evaluation Method

Demonstration

Joint Staff Operational Criteria Tested

2 (2.1, 2.1.1, 2.1.2, 2.1.3, 2.1.4, 2.3, 2.3.1, 2.3.2, 2.3.3, 2.3.4)
3 (3.1, 3.1.1, 3.2, 3.2.1, 3.3, 3.3.1)
5 (5.1, 5.1.1, 5.1.2)

Configuration

Tests were executed in a closed network environment using automated test tools on TDM network devices (Promina 100 and Cisco 3700 series routers). For interoperability testing, the Ixia Chariot generated concurrent, multi-protocol traffic (data, voice, and video) and report results for IPv4 and IPv6.

Performance tests were run using the Spirent SmartFlow software on the SMB-600C SmartBits chassis to evaluate the compliance of a system or component with Joint Staff IPv6 operational criteria 3 performance requirements. The SmartFlow automated tool set tested the performance of the DUTs (Promina 100 and Cisco 3700) in a closed, controlled environment from a packet-loading standpoint.

Results//Lessons Learned

Interoperability

Protocol-based interoperability test results were split into two categories dependent upon script function: transaction-based and streaming-based. The following scripts were transaction based:

- FTP (Get/Put)
- HTTP (Text/Graphics Interchange Format)

- HTTPS
- POP3
- SMTP
- SNMP
- Lightweight Directory Access Protocol (LDAP)
- SIP

Transaction based scripts resulted in 100% successful transport and delivery for all protocols and traffic types tested. Because of the immaturity of the current IPv6 technology, the protocols listed below could not be successfully tested in a pure IPv6 environment.

- Resource Reservation Protocol
- RTP

Streaming-based scripts successfully completed when running concurrently with transaction-based scripts with no significant loss. Aggregate streaming results for all tests resulted in a less than 3% difference between IPv4 and IPv6 with respect to packet loss.

- DNS
- G.711u (VoIP)
- IPTV – Video
- IPTV - Audio

Performance

Seven bandwidths were tested, as four performed with greater latency when using IPv6. The average latency for these four bandwidths was 1.36% longer using IPv6 than when using IPv4. Three of the seven bandwidths tested performed with less latency when using IPv6. These three bandwidths averaged .19% less latency when using IPv6 than when using IPv4. The operational impact of these small differences in latency is seen as inconsequential. Table D-14 lists detailed performance results.

Table D-14 Low Bandwidth Performance Results

Bandwidth Kb/sec	Average Latency IPv4 in µs	Average Latency IPv6 in µs	Percent IPv6 compared to IPv4	Interrelated Latency Percent
8	12901757.5	13006776.8	99.19%	.81% More latency
16	15135184.8	15506218.6	97.61%	2.39% More Latency
32	2615812.15	2614444.51	100.05%	.05% Less Latency
64	8549650.51	8511554.67	100.45%	.45% Less Latency
128	766341.364	765736.597	100.08%	.08% Less Latency
256	3126812.15	3131081.47	99.86%	.14% More Latency
512	1243633.19	1270136.92	97.91%	2.19% More Latency

Conclusions

Interoperability

Transaction based scripts resulted in 100% successful transport and delivery for all protocols and traffic types tested. IPv6 transmitted streaming based scripts at least 97% as well as IPv4 when transmitting concurrently with transaction-based scripts establishing interoperability for the protocols tested.

Performance

Latency tests executed on the network under test resulted in equivalent performance when comparing IPv4 to IPv6. Less than 4% variance in the packet transmission rates was recorded between the protocols. The test showed IPv4 and IPv6 having equivalent performance.

D.20 IPv6 Autoconfiguration White Paper

Testing Organization and Publication Date

SPAWAR Systems Center
5 February 2007

Summary

The Navy is leading a joint effort among the Services to develop the Joint Tactical Edge Networks (JTEN) concept for providing connectivity to highly mobile and disadvantaged users at the tactical edge network where fixed infrastructure may not be supported. This paper compares the current mechanisms available for IPv6 autoconfiguration and recommends an efficient address autoconfiguration mechanism that should make it easier for DoD/tactical network deployment and transition to IPv6.

Test and Evaluation Method

Engineering Analysis

Joint Staff Operational Criteria Tested

10 (10.2, 10.2.1)

Configuration

The JTEN includes a hierarchy of MANET that enable communicating nodes to self-organize their own network automatically and rapidly. The goal of the JTEN concept is to help warfighters shorten the kill chain in a battlespace scenario.

One technique to shorten the kill chain is the plug-and-play capability for nodes to obtain and configure valid IPv6 addresses automatically and quickly. An efficient address autoconfiguration mechanism is needed to cope with the highly dynamic nature of MANET. The autoconfiguration mechanism should allow nodes to self-organize quickly and efficiently in order to reduce the time needed to deploy tactical networks and to remove the burden and inflexibility of manual configuration where the exact network condition is unpredictable.

Two IPv6 address autoconfiguration mechanisms can be implemented to support the JTEN operation concept. The stateless approach, based on the Neighbor Discovery (ND) protocol, enables hosts to automatically configure their own IPv6 addresses without a server. The stateful approach, based on DHCPv6, requires a DHCP server to provide hosts (clients) IPv6 addresses and other information such as DNS. However, both approaches have more difficulties in a MANET environment than in wired networks. This is due to instability of links, multi-hop topology that is highly dynamic because of frequent partitioning, merging of network nodes, and the absence of central administration. The stateful, DHCP approach is too centralized and impractical, as it requires a server, which may not be reachable if MANET partitions. The stateless approach uses the combination of network prefix and the host's 64-bit "interface identifier" to configure an IPv6 address that is unique within a controlled local MANET. Thus, these addresses are valid and routable within a MANET only, but are not unique and routable globally.

Results

In supporting a JTEN requirement to enable MANET at the tactical edge network connecting with the GIG, the gateways approach is an extension of standard stateless autoconfiguration in order to provide the following autoconfiguration capabilities for MANET:

- Guarantees the uniqueness of IPv6 address assignments and is routable globally
- The ability to cope with the network dynamics present in MANET
- Scalable (e.g., thousands of nodes with multiple interfaces)
- Independent of routing protocols.

The gateways approach can be implemented in extend to the standard stateless autoconfiguration in order to provide ad-hoc nodes globally routable address assignments and connectivity to the fixed network infrastructure such as the GIG. This mechanism enables MANET nodes to build valid global IPv6 addresses when the MANET is interconnected to an external, fixed network infrastructure (e.g., the GIG) through one or more gateways. Nodes that are not directly linked to the GIG can use multi-hop paths to reach the gateways that forward outbound traffic to the host.

Conclusions

In an effort to support the JTEN concept, the gateways approach can be implemented to the standard stateless autoconfiguration in order to cope with the multi-hop topology, and to automatically assign globally unique and routable IP addresses to nodes in a MANET for connecting to the GIG. This new approach can benefit the JTEN concept by helping warfighters shorten the kill chain in a battlespace scenario. This autoconfiguration mechanism should allow tactical nodes to self-organize quickly and efficiently in order to reduce the time needed to deploy tactical networks and remove the burden and inflexibility of manual configuration where the exact network condition is unpredictable.

D.21 Operational Issues with IPv6 DNS

Testing Organization and Publication Date

SPAWAR Systems Center
28 January 2007

Summary

This report discusses the ongoing operational issues and shortcomings related to DNS with IPv6, and recommends actions that should be taken in response to these issues and shortcomings.

Test and Evaluation Method

Engineering Analysis

Joint Staff Operational Criteria Tested

2 (2.1, 2.1.1, 2.2, 2.2.1)
8 (8.1, 8.1.1)

Configuration

A typical network is configured and functions as follows: a DNS client, also known as stub resolver, issues a query for host name "host1.example.com" to a caching/local DNS server, also known as resolver, which then handles the entire name resolution recursively. The local DNS server first sends the query to the root name server (NS) which is statically configured on the local DNS server. The root server answers the local server and advises it to contact the NS ".com." The local DNS server sends the same query to the NS ".com," which answers and advises it to contact the name server "example.com." The local server repeats the above steps in recursive manner until it receives the IP address for host1.example.com from the name server "example.com". The local DNS server finally caches the result and returns it to the client.

Results/Issues

A summary of some of the major issues with IPv6 DNS implementation are:

- Standards: AAAA versus A6
    - AAAA records are preferable at the moment for production deployment of IPv6
    - A6 records have interesting properties that need to be better understood before deployment
    - It is not known if the benefits of A6 outweigh the costs and risks
    - RFC 3363 recommends AAAA records should be used in preference of A6 records for deployment of IPv6.

- Stateless Autoconfiguration
  - There is no standard method for stateless autoconfiguration of IPv6 hosts to discover a DNS server's address; instead, DHCPv6 or manual configuration must to be used. Currently, the DHCPv6 approach is preferred
  - There is no method yet for populating reverse DNS data for a network, particularly for stateless autoconfigured hosts. Such reverse lookups are used as a weak authentication in some instances, e.g., by sendmail to accept SMTP from local hosts.

- IPv6 or IPv4 Transport
  - When a dual-stack host looks up a dual-stack destination host that has both A and AAAA records in the DNS server and does not have (global) IPv6 connectivity, it typically first tries querying the DNS for an AAAA record using IPv6 transport in preference to IPv4 transport. It turns out that an IPv6 packet is sent but there is no reply. After some timeout, a fallback to IPv4 occurs, but introduces unnecessary delay
  - If a client receives both AAAA and A records for the same destination from DNS queries, it typically connects to the IPv6 service first and only defaults to IPv4 if IPv6 fails. In some circumstances, the address selection prefers AAAA records to A records, even when the destination node does not have IPv6 connectivity. After some timeout, a fallback to IPv4 should happen, which works, but introduces unnecessary delay.

To avoid the above issues, AAAA records of a node should not be added to the DNS until all services of the node are IPv6-enabled and have IPv6 (global) connectivity.

Conclusions

A dual-stack-capable DNS hierarchical system should be in place prior to implementation of dual stack on DoD network infrastructure for transition to IPv6. Ongoing operational issues of DNS IPv6 must be identified and resolved for successful deployment of IPv6.

D.22 IPv6 Multihoming White Paper

Testing Organization and Publication Date

SPAWAR Systems Center
30 January 2006

Summary

This report discusses possible multihoming approaches in an effort to support the Navy architecture and addressing plan for successful transition to IPv6. The goal was to identify solutions for providing load-balancing, redundancy, and traffic engineering in order to cope with the bandwidth constraints issue presented in the Navy architecture.

Test and Evaluation Method

Engineering Analysis

Joint Staff Operational Criteria Tested

10 (10.1, 10.1.2)

Configuration

There are two possible multihoming solutions for IPv6: Routing and Host-Centric. Some routing approaches (e.g., BGP and Cross-tunnels at Site Exit Routers) are mature but may not be practical for Navy/tactical networks. Some Host-centric approaches (e.g., shim Layer and Name, Address and Route System) though practical and beneficial to Navy/tactical networks, need further research, testing and development due to their immaturity.

Results

In Routing approaches, hosts use a single IPv6 address and the mechanisms for supporting IPv6 multihoming are handled by routers. The routing approach can implement the cross-tunnels at site exit routers in order to allow multiple prefixes inside a multihoming site. This approach relies on the use of tunnels in order to provide multiple prefixes and route aggregation; therefore, it introduces complexity when routes are withdrawn and requires renumbering in the event of primary link failure. This approach introduces a scalability problem in the Internet routing table, incomplete route summarization, and inflexibility due to the use of a single prefix in the site.

Host-Centric IPv6 multihoming provides solutions by using multiple prefixes and providing fault-tolerance, route aggregation and traffic engineering. The Host-Centric approach can support the current Navy IPv6 addressing scheme that provides maximum possible route summarization necessary to prevent IPv6 route explosion on Navy networks, and minimizes the impact of network overhead on low-bandwidth RF tactical links.

MIPv6 was often proposed as a multihoming solution in the early stages of development for IPv6 multihoming, since the preservation of established communications through movement is similar to the preservation of established communications through outages in multihomed sites. MIPv6 also introduces return routability, but the drawback of return routability is the reliance on the availability of the HA; therefore, MIPv6 is not a suitable mechanism for IPv6 multihoming.

Conclusions

In conclusion, the Host-Centric approach is recommended for supporting the current Navy addressing scheme. In addition, the Efficient Security for Multihoming architecture can be implemented to secure ship-to-shore communications from redirection attacks. These capabilities provide security and improve ship communications performance over bandwidth-constrained satellite communication links. However, these approaches are not completely mature and need further research, development, and testing with current IPv6 technologies to meet specific needs of the DoD/Navy.

D.23 Special Interoperability Test Certification of Microsoft Windows Vista Enterprise Operating System installed on a Panasonic Toughbook CF-74 and a Panasonic Toughbook CF-51 for Internet Protocol Version 6 (IPv6) Capability

Testing Organization and Publication Date

JITC
March 2007

Summary

This report presents the results of the Special Interoperability Test Certification of Microsoft Windows Vista Enterprise OS installed on Panasonic Toughbooks CF-74 and CF-51. Testing occurred from 15 January to 16 February 2007 at JITC's AIPTL. Upon successful testing, the OS and Personal Computers (PC) were placed on the DoD APL.

Test and Evaluation Method

Exercise

Joint Staff Operational Criteria Tested

1 (1.1, 1.1.1, 1.2., 1.2.1)
2 (2.3, 2.3.1)
8 (8.1, 8.1.1)

Configuration

The OS and PCs were tested as part of a simulated DISN Core node test architecture managed by the AIPTL. The core consisted of hardware and software listed below.

Table D-15 Equipment Configuration

Equipment Name		Model Number	IOS/OS Version(s)
Hardware	Cisco Router	CISCO3845	12.3(14)T2
	Cisco Router	CISCO3845	12.4(4)T1
	2 Juniper Routers	Juniper M40e	V 7.6R3
	2 Juniper Routers	Juniper T320	V 7.6R3
	Juniper Router	Juniper T640	V 7.6R3
	6 Dell Power Edge Servers	2850	Microsoft Windows 2003
	Panasonic Tough Book	CF-51	Microsoft Windows Vista Enterprise
	Panasonic Tough Book	CF-74	Microsoft Windows Vista Enterprise
Software	Wireshark	N/A	Enterprise
	Dibbler	N/A	V.0.99.2
	Cisco Router	N/A	0.4.1

Results

The following table summarizes the applicable Joint Staff IPv6 operational criteria results from the Microsoft Windows Vista Enterprise OS on the Panasonic Toughbooks CF-74 and CF-51 certification test.

Table D-16 Test Results

RFC Title	Testing Completed		Host/Workstation	
	Conformance	Interoperability	Requirement	Met/Not Met
Privacy Extensions for Stateless Address Autoconfiguration in IPv6	Stated in LoC	Yes	Yes	Met
Default Address Selection for IPv6	Stated in LoC	Yes	Yes	Met
DNS Extensions to Support IPv6	Stated in LoC	Yes	Yes	Met
Uniform Resource Identifier (URI): Generic Syntax	Stated in LoC	Yes	Yes	Met
File Transfer Protocol	Stated in LoC	Yes	Yes	Met
Simple Mail Transfer Protocol	Stated in LoC	Yes	Yes	Met
Hypertext Transfer Protocol		Yes	Yes	Met
IP Authentication Header	Stated in LoC	Yes	Yes	Met
IP Encapsulating Security Payload (ESP)	Stated in LoC	Yes	Yes	Met
Cryptographic Algorithm Implementation Requirements for ESP and Authentication Header (AH)	Stated in LoC	Yes	Yes	Met
Transition Mechanisms for IPv6 Host and Routers	Stated in LoC	Yes	Yes	Met

Conclusions

Microsoft Windows Vista Enterprise OS, installed on the Panasonic Toughbooks CF-74 and CF-51, met the IPv6-capable requirements and are certified for listing on the DoD APL as an IPv6-capable Host/Workstation.

D.24 Special Interoperability Test Certification of Techguard PoliWall for Internet Protocol Version 6 (IPv6) Capability

Testing Organization and Publication Date

JITC
March 2007

Summary

This special certification is based on IPv6-capable testing conducted by JITC at Fort Huachuca, Arizona. Testing was conducted at JITC's AIPTL from 19 to 23 February 2007. The PoliWall provides simplified network security and communications traffic management by acting as a transparent bridge between a firewall and an external network. Upon successful testing, the device is certified for listing on the DoD APL as IPv6 capable as a Network Appliance.

Test and Evaluation Method

Exercise

Joint Staff Operational Criteria Tested

1 (1.1, 1.1.1, 1.2., 1.2.1, 1.4, 1.4.1)
8 (8.1, 8.1.1)

Configuration

The operational architecture was the simulated Network Boundary composed of the equipment listed below.

Table D-17 Equipment Configuration

Equipment Name		Model Number	IOS/OS Version(s)
Hardware	Cisco Router	CISCO3745	12.3(14)T2
	Cisco Router	CISCO3745	12.4(4)T1
	Dell Power Edge Server	2850	Microsoft Windows Server 2003
	Dell Power Edge Server	2650	Microsoft Windows Vista Enterprise
	Dell Power Edge Server	4600	RedHat Enterprise 4
	Gateway Laptop	Gateway Laptop	Microsoft Windows XP
	PoliWall	PoliWall	1.20.40
Software	Microsoft Windows Vista	N/A	Enterprise
	Microsoft Windows Server	N/A	2003
	RedHat Enterprise	N/A	4
	Microsoft Windows XP	N/A	Professional
	Wireshark	N/A	V.0.99.2
	PoliWall	N/A	1.20.40

Results

The Techguard PoliWall effectively provided network security and managed associated traffic. Table D-18 lists the applicable Joint Staff IPv6 operational criteria that were assessed.

Table D-18 Test Results

RFC Title	Testing Completed		Network Appliance	
	Conformance	Interoperability	Requirement	Met/Not Met
IP Authentication Header	Stated in LoC	Yes	Yes	Met
IP Encapsulating Security Payload (ESP)	Stated in LoC	Yes	Yes	Met
Cryptographic Algorithm Implementation Requirements for ESP and Authentication Header (AH)	Stated in LoC	Yes	Yes	Met
Cryptographic Suites for IPSec	Stated in LoC	Yes	Yes	Met
Transition Mechanisms for IPv6 Host and Routers	Stated in LoC	Yes	Yes	Met

Conclusions

The Techguard PoliWall met all the test requirements of a network appliance and is certified for listing on the DoD APL as an IPv6-capable Network Appliance.

D.25 IPv6 Protocol Security Assessment and Issues

Testing Organization and Publication Date

DITO/SI International
15 March 2007

Summary

This report provides a high-level protocol security assessment of issues associated with the migration to IPv6 by DoD organizations. This document may also be used to develop and execute IPv6 IA and T&E plans.

Test and Evaluation Method

Engineering Analysis

Joint Staff Operational Criteria Tested

1 (1.1, 1.1.1, 1.2, 1.2.1, 1.4, 1.4.1, 1.4.2, 1.5, 1.5.1)
7 (7.1, 7.1.2)
8 (8.1, 8.1.1, 8.1.3)
9 (9.1, 9.1.1, 9.2, 9.2.1)
10 (10.2, 10.2.1)

Configuration

While IPv6 security assessment issues may arise in several IPv6 transition areas; such as IPv6 protocols and services, IPv6 protocol implementations, applications considerations, network architecture, and co-dependent technologies, this document covers only the IPv6 protocol and services.

Results

Mobility

The following types of security issues can arise from IPv6-based mobility:

- Easy ability to determine address spaces, e.g., MAC-layer identifications
- ICMPv6 discovery/resolution/redirect services may allow access leading to system processor loading and DoS
- IPv6 mobile Binding Update authorization allows end users to force or spoof Bus leading to traffic interception and rerouting.

Some mobility security issues can be mitigated using stringent ACL and firewall filtering techniques. Other issues can be fixed using strict authentication techniques, such as IPSec AH.

Transition Mechanisms

Dual stack presents at least two security issues:

- IPv6 DNS security issues
- DHCPv6 security issues.

Implementing tunnels (IPv6 over IPv4, IPv6 to IPv4, Intra-Site Automatic Tunnel Addressing Protocol, and GRE) presents the following security and administrative issues:

- Manual tunneling has a high administrative overhead, but is recommended by the IETF Security Working Group
- Automatic tunneling is preferred if a large number of tunnels are required, but there is no mechanism to authorize the tunnel creator
- IPv6 multicast security issues
- Unauthorized router access
- Access and security between tunnel end-points can be compromised.

Some of these issues may be resolved as designs and implementations mature, including software applications and firmware on the network device line-cards.

Network Management and Operations (NM/OPS)

Overlapping IPv4/IPv6 NM/OPS functions caused by configuration errors and bad code may allow access leading to traffic interception, rerouting, DoS, addressing spoofing and ingress filter avoidance. These issues include:

- System access and configuration
- System default routers can be changed/altered
- User ID may not be known and user may not be authenticated
- IPv4/IPv6 internetworking interaction is not well defined in the IETF; access to one side of the network allows access to the other side of the network.

Conclusions

The introduction of IPv6 into DoD operational networks may pose security risks and introduce potential vulnerabilities in addition to those inherent in current IPv4 networks. While existing DoD vulnerability assessment techniques should provide a level of security equivalent to that of IPv4 networks, the issues regarding IPv6 security assessments must continue to be analyzed.

D.26 JCS Criteria 4, Phase 1 – Demonstration of QoS Capabilities of IPv6 Using DiffServ (FY07 Moonv6 Demonstration)

Testing Organization and Publication Date

SPAWAR
23 December 2006

Summary

This document describes tests that demonstrated IPv6 QoS capability using DiffServ during the Moonv6 demonstration for FY 2007. The demonstration was carried out between the SPAWARSYSCEN and at other participant test beds. Laboratory testing of DiffServ in a limited-bandwidth environment conducted within the SPAWAR test center is presented within this report.

Test and Evaluation Method

Exercise

Joint Staff Operational Criteria Tested

4 (4.1, 4.1.1)
5 (5.1, 5.1.1)
8 (8.1, 8.1.1)

Configuration

Moonv6 Demonstration

A Cisco 3625 router was connected to the DREN via the SPAWARSYSCEN campus network. Static routing was used in the test bed to connect to the campus network and the campus network advertises routes into the DREN using the BGP routing protocol.

The hosts were connected to the test router, a PC hosting an IxChariot Endpoint, a SmartBits 600B Performance Analysis System and a laptop PC hosting a network sniffer (Ethereal). All addresses contained 64-bit prefixes. Data traffic was generated by the IxChariot Endpoint or SmartBits System. DiffServ markings were made by either the IxChariot Endpoint or SmartBits System when the traffic was generated. The packets could also be tagged by the router according to a policy configured in the router.

Remote sites consisted of JITC, Fort Huachuca, Arizona, Air Force Communications Agency and other DoD organizations. The JITC test bed had a group of routers that emulated the GIG core and its DiffServ policy.

The IxChariot Endpoint is also attached to the GIG core via an Edge Router that marked packets with DiffServ Code Point (DSCP) according to the DSCP assignments proposed for the GIG.

Laboratory Testing

The laboratory testing setup was similar to the Moonv6 test setup, with the exception that an ADTech Model SX/12 Data channel simulator was substituted for the WAN part of the network, the DREN, and simulated GIG core at JTIC. Two sites were connected via the ADTech Satellite Simulator (SATSIM). The SATSIM was configured for a throughput of 256 Kbps.

Results

Moonv6 Demonstration

The DSCP markings were applied to packets passed between SPAWARSYSCEN in San Diego and JITC. Traffic DSCP marking was tested using IxChariot Endpoints and SmartBits. Using IxChariot Endpoints, markings were generated and transmitted from the lab at SPAWAR to JITC. All packets were marked and received with the same markings. DSCP markings were successfully passed via the simulated GIG core network at JTIC.

Laboratory Testing

<u>Separate Flooding of Voice, Video, and Data Traffic Classes</u>

Table D-19 summarizes a test of overloading separately the traffic of each type: “voice,” “video,” and “Other TCP.” This represents three separate tests that each queue was guaranteeing for the bandwidth allocated to the queue.

Table D-19 Separate Flooding of Voice, Video, and Data Results

Queue	Traffic Type	Reserved (Kbps)	Offered (Kbps)	Allocated (Kbps)	Received (Kbps)
2	CM-Voice	76	96	76	76
1	Video	76	100	76	96
0	Other TCP	64	135	64	96

Voice received exactly the bandwidth allocated because it was priority-queued. This was expected since voice is a priority queue and voice packets above the allocated rate were dropped. Video and other TCP received more bandwidth than was allocated to those two queues. This was also expected because both queues used Class Based Weighted Fair Queuing and is, by design, able to utilize bandwidth available from other queues that are not fully utilized.

<u>Simultaneous Flooding of Voice, Video and Data Traffic Classes</u>

Traffic of each of the 4 classes was offered to the link at rates greater than the allocated rates and the received rates were measured. Table D-20 presents the results.

Table D-20 Simultaneous Flooding of Voice, Video, and Data Results

Queue	Traffic Type	Reserved (Kbps)	Offered (Kbps)	Allocated (Kbps)	Received (Kbps)
3	Control	38	38	38	38
2	CM-Voice	76	89	76	76
1	Video	76	134	76	75
0	Other TCP	64	67	64	64

All four queues delivered the approximate allocated rates configured.

Conclusions

The basic QoS functionality using DiffServ mechanism with a commercial Cisco router was tested. The basic features of QoS for IPv6 were tested and found to be supported and functional for the DiffServ mechanism of providing QoS for different classes of traffic (voice, video, and data).

One discrepancy found is that Low Latency Queuing is not supported yet by Cisco for IPv6 by the Cisco IOS version used for this testing, IOS 12.4. As a work-around, a priority queue was configured for the voice traffic, which provided approximately the same functionality.

D.27 Test Results and Lessons Learned

Testing Organization and Publication Date

Air Force System Networking (AFSN)

Summary

The Air Force was tasked to test an IPv6 addressing scheme, test IPv6 routing protocols, and test IPv6 tunneling technologies in an Air Force simulated network environment. After testing, performance results, discrepancies, recommendations, and lessons learned were provided in this report.

Test and Evaluation Method

Demonstration

Joint Staff Operational Criteria Tested

1 (1.1, 1.1.1, 1.2, 1.2.1)
3 (3.1, 3.1.1, 3.3, 3.3.1)
8 (8.1, 8.1.1, 8.1.2)

Configuration

The AFSN test and integration facility was set up to simulate the DISA WAN and three individual mock bases (Tyndall, MacDill, and Eglin). Equipment typically found in the Air Force network was utilized (Cisco 2950 and 3560 switches and Cisco 7206 VXR, 3725 and 3745 routers). The DISA WAN consisted of three core 7206 VXR routers (core1, core2, and core3). The network at Eglin was set up with dual paths and dual Shelf Discovery Protocol (SDP) routers to simulate future Combat Information Transport System (CITS) Block 30 architectures. The network at Tyndall and MacDill had single SDPs and single paths to the WAN. Each base had an external router below the SDP router and switches between all routers. All connections between routers and switches were at 100 Mbps. The DISA core routers were set up to run BGP between all core routers and the SDP routers at each base. The routers in Tyndall were set up to run RIP within the base network. The network at Eglin ran Enhance Interior Gateway Routing Protocol within the base network. The routers located in MacDill ran OSPF within the base network Workstations and servers were placed at various endpoints. Spirent test equipment was used to load the test network and recorded performance results.

Results

This report covers results from 22 separate scenarios. Below are examples of some of the tests that were run and their results.

100% IPv4 Traffic over Stacked Configuration

Results were very similar to the baseline testing on the all IPv4 network tested. However, while no or minimal losses were again seen up to about 81% load, at this point, some higher losses were noted compared to the initial baseline. Running 100% IPv4 over the stacked network configuration caused little or no additional loss compared to the original baseline test for frame sizes over 256 bytes.

50% IPv6 Traffic Over Stacked Configuration

Multiple traffic flows were used across the network. Increasing traffic to 50% IPv6 produced only minimal increases in loss. Some loss was noticed at 46% loading for the 128 byte frame size. The increase in loss was minimal and no increase in loss was seen in larger frame sizes.

100% IPv6 Traffic Over Stacked Configuration

Multiple IPv6 traffic flows were used across the network. As would be expected with all IPv6 traffic, throughput and frame loss numbers increased slightly more compared to the 75% test conducted earlier. At 128 byte frames with 91-96% loading, one fourth of the traffic paths experienced 100% loss. However, as frame sizes increased, loss decreased. As was the case in all the other tests, once frame sizes of 512 bytes were reached, no more significant loss was experienced. Even at smaller frame sizes, losses were only significant with larger loads.

100% IPv6 Traffic Over IPv6 Configuration

Multiple IPv6 traffic flows were used across the network. There was less frame loss and throughput loss than was experienced in testing over a dual-stacked network. No losses were seen until loading for 128 byte frame sizes reached greater than 51% (compared to losses beginning at 41-46% loading for the stacked configuration). At no time did any paths experience 100% loss, no matter how high the loading. Again, with larger frame sizes, no loss was seen. These results are consistent with the knowledge that running an IPv6 only configuration should have lower processor utilization than running a dual-stacked IPv4 and IPv6 configuration.

Tunneling IPv4 Over IPv6 – GRE Tunneling Over IPv6 (GRE ipv6)

This test tunneled one flow of IPv4 traffic over a network consisting of equipment running only IPv6 (addresses on interfaces, routing protocol stacks, etc.). The tunnel was conducted between the DISA core and Eglin as the two endpoints. The routers at the endpoints were configured as GRE tunnels. The results indicated extremely high throughput and frame loss. After further discussions with the vendor, it was again confirmed that Cisco routers do not support Cisco Express Forwarding (CEF) for IPv4 tunneling over IPv6 networks. Therefore, processor switching is utilized and any router performing this type of tunnel will experience high loss and poor performance for that data. However, all the losses were just a few percentage points higher than the generic tunneling test. GRE tunnels performed just slightly worse than generic tunneling over IPv6.

Tunneling IPv6 Over IPv4 – GRE Tunneling Over IPv4

This test tunneled multiple flows of IPv6 traffic over a network consisting of equipment running only IPv4. The tunnels were conducted between the DISA core and Eglin as the two endpoints. The results indicated relatively high throughput and frame loss. The added traffic caused by the extra flows forced this tunnel to perform poorly. In addition, as the test for each frame size neared 100% loading, it was noticed that losses came very close to 100% (unlike the single flow test, which did not get as close to reaching 100% loss).

Tunneling IPv6 Over IPv4 – Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)

This test tunneled one flow of IPv6 traffic over a network consisting of equipment running only IPv4. The tunnel was conducted between the DISA core and Eglin as the two endpoints. The routers at the endpoints were configured as ISATAP tunnels. The results were almost identical to the 6to4 tests for single flow, with just a slightly higher loss on maximum loaded smaller frame sizes. With medium to larger frame sizes, the results were practically identical to the 6to4 test for all loading. ISATAP tunnels need no tunnel destination address in the configuration (though endpoints do need to appear in the routing tables). They are easily configured when compared to manual tunnels.

Conclusions

Comparing preliminary baseline results with IPv6 configuration test results reveals that there is additional throughput and frame loss on systems processing IPv6 packets. Tunneling tests revealed that many tunnel types are based on using processor switching instead of CEF. This method is very processor intensive and poor performance can be anticipated with any tunnel method that uses processor switching. Tunnel types that require processor switching (instead of CEF) should only be used as a last resort. Specifically, refrain from using tunnels over IPv6 networks and GRE tunnels on IPv4 networks. Using Cisco IOS 12.4 or later to permit implementation of all IPv6 processes and features is recommended. Larger frame sizes should be used when possible. Larger frame sizes are more efficient (with less overhead) and, therefore, produce better performance and higher throughput. On dual-stack transition testing, almost no significant losses were ever experienced on frame sizes larger than 512 bytes. High-end routers should be used where possible; avoid using low-end routers that have higher processor utilization, and therefore less throughput and more frame loss. Manual configuration of all network equipment should be performed, although autoconfiguration and DHCPv6 may have usefulness on hosts and servers.

D.28 Network Management IPv6 Feasibility Study Report

Testing Organization and Publication Date

Air Force Research Laboratory / Northrop Grumman
31 January 2007

Summary

The objective of this effort was to provide analysis, design, development, integration, and testing in support of demonstrating the ability to move network elements to other locations while maintaining connectivity via their original IPv6 addresses using MIPv6 within the Joint Capability for Airborne Networking (JCAN) system. This report summarizes the deficiencies of IPv4 and the advantages of IPv6 that will enable future 21st Net-Centricity.

Test and Evaluation Method

Engineering Analysis

Joint Staff Operational Criteria Tested

7 (7.1, 7.1.1, 7.1.2, 7.1.3)

Configuration

The JCAN system architecture consists of three major elements: an Airborne MN, a Ground Node (GN), and one or more Ground Entry Sites (GESs).

The JCAN MN computers manage data routing, application services, data logging, and the user interface for JCAN system monitoring and control. The MN provides the interface between the E-8C LAN and existing E-8C Ultra High Frequency (UHF) and High Frequency (HF) aircraft radios. The connection between the MN and the E-8C LAN is accomplished via a standard Ethernet connection. The aircraft LAN supports 18 operator workstations. The connection between the MN and the existing E-8C radio equipment is accomplished via UHF and HF Radio Interface Modules (RIMs). The RIMs provide the conversion from a standard Ethernet connection on one-side and radio specific serial interfaces for the UHF or HF radios on the other side. Through digital control from the MN, the RIMs provide the capability to switch between standard E-8C voice radio operation and JCAN data operation. JCAN operation via the existing HF radio links also requires the addition of HF modems and KIV-7 crypto units for each radio. UHF operation uses existing KY-58 crypto units to provide secure data operation.

The JCAN GN is similar to the JCAN MN, interfacing with the JCAN enabled radios at one or more GESs. The GESs can be collocated with the JCAN GN or geographically separated. This deployment supports distributed GESs. Each GES is configured with multiple AN/URC-200 radios, KY-58 crypto units, UHF Tactical Communications antennas and a JCAN Serial

Interface to Military Radios. The GESs are connected to the JCAN GN through a satellite interface.

JCAN extends MIP capabilities to support mobile networks, secure registration, concurrent multi-path routing over multiple heterogeneous links, mission-based QoS, and Performance Enhancing Proxies.

Results

The following are the features of the IPv6 protocol that are advantageous for JCAN capability:

- Large Address Space (simplifies management and administration of the JCAN system)
- Efficient and Hierarchical Addressing and Routing Infrastructure
- Stateless and Stateful Address Configuration
- Built-in Security
- New Header Format – Extensibility
- Better Support for QoS (greater ability to support QoS differentiation)
- Enhanced Neighbor Discovery Mechanism (IPv6 enforces topologic consistency)
- No Foreign Agent Deployed.

The following are concerns about using IPv6 in JCAN:

- Decreased Transport Efficiency (increased IP Header size)
- Further Decreased Transport Efficiency (lack of mobility support mode within the MIPv6 standard forces the use of tunnels that originate and terminate at the mobile node).

Conclusions

IPv6 does not by itself provide a solution to airborne networking problems. However, the available address space alone will open up some opportunities for the implementation and deployment of the airborne network.

In the remainder of this effort, a MIPv6 version of JCAN will be developed and implemented then explored, compared and contrasted to realize any benefits of this implementation over the current MIPv4 solution.

D.29 IPv6 Vulnerability Assessment Report for the Air Force Standard Desktop Configuration of Microsoft Windows Vista

Testing Organization and Publication Date

Air Force Information Operations Center (AFIOC)
20 April 2007

Summary

The Air Force IPv6 Team investigated the effect Windows Vista would have on an emulated Air Force network. This test was conducted in order to provide confidence that IPv6-enabled nodes would not compromise the security of the operational enterprise network. This report consists of examinations designed to ensure Windows Vista provides a high level of security, even when IPv6 is enabled. The areas evaluated include accessibility from the network; operability of applications and services; and related features that may have significant weaknesses based on known or unknown criteria of use.

Test and Evaluation Method

Experiment

Joint Staff Operational Criteria Tested

1 (1.4, 1.4.1, 1.4.2, 1.5, 1.5.1)
8 (8.1, 8.1.1)

Configuration

Testing was conducted in a dual-stack configuration that emulated an Air Force network. The equipment used in the test included:

- Cisco 3825 Routers
- Cisco 7206 Router
- Cisco Switches
- Computers loaded with
    - Vista Enterprise v2.0
    - XP Service Pack (SP) 2
    - Longhorn Beta 1
    - Server 2003 SP 1
- Hewlett Packard Printers
- MU-4000 Analyzer
- Ixia Traffic Generator.

Results

Some of the interesting findings from this assessment were either expected or validated from previous beta versions of Windows Vista. One of the outstanding findings was the interdependencies of protocols seen when trying to enumerate ports and services on the system. This shows how Windows Vista has changed its implementation of the Remote Desktop Protocol (RDP) connection and its use of TCP port 443, which is usually reserved for HTTP over Secure Sockets Layer. The current configuration is still only using TCP port 3389 or RDP.

Significant findings in the Windows Firewall implementation shows it was able to block all ICMPv6 message types, but did not have the granularity to block an unknown data-link layer address of all zeros. Blocking all ICMPv6 message types is not feasible because certain messages are used for the discovery of other hosts on the network and routing capabilities to other subnets, which is crucial to the operation of IPv6 on the network.

Other findings showed that Netcast with IPv6 support could also be used to enumerate or fingerprint the Windows Vista OS when certain ports are open. However, when the Windows Firewall is invoked and services blocked, no ability to enumerate the system could be established. Specific examinations that failed were related to fragmentation of packets and the ability to reassemble a complete packet only when the fragment identification value is the same as the original packet.

Conclusions

This assessment has two clear conclusions regarding the IPv6 implementation in Windows Vista: 1) The IPv6 implementation in Windows Vista has improved weaknesses found in previous Windows OSs, and 2) To evaluate the full scope of IPv6 and the expanding protocols surrounding it, additional assessments with more rigorous security tools should be conducted.

The security areas evaluated in this assessment are predominately a baseline for all evaluators to use on the Windows Vista OS. Continued techniques or methods will be introduced to evaluate Windows Vista, but collaboration is essential to finding any significant weaknesses in the new OS.

D.30 IPv6 Scalability Testing Final Report

Testing Organization and Publication Date

DISA
30 March 2007

Summary

This document presents the results of the System Level testing executed in order to verify the scalability of IPv6 in an operational environment. The objective was to develop technical data and information on the Joint Staff Operational Criterion 6, scalability of IPv6 in the DoD environment, so that the Chairman of the Joint Chiefs can certify that the conversion of DoD networks to IPv6 will provide equivalent or better scalability than the current IPv4.

Test and Evaluation Method

Experiment

Joint Staff Operational Criteria Tested

6 (6.1, 6.1.1, 6.1.2, 6.1.3, 6.1.4)
8 (8.1, 8.1.1)

Configuration

The testing was divided into a series of functional areas – routing, subnets, multicast, and services. The routing test was designed to demonstrate that IPv4 and IPv6 routing protocols can scale to the same levels. The subnet testing was designed to show that IPv4 and IPv6 subnets can handle the same number of hosts. The multicast test section was set up to show that IPv4 and IPv6 can support the same levels of multicast service. The services section was designed to show that basic network services such FTP and HTTP can handle both IPv4 and IPv6 requests.

Traffic parameters were based on assumptions of how IPv6 would be integrated into the network. The assumption was that IPv6 integration would be a slow and steady process. The following traffic profiles were used as evaluation points for the integration process:

- 100% IPv4
- 90% IPv4/10% IPv6
- 50% IPv4/50% IPv6
- 10% IPv4/90% IPv6.

Packet size was another important parameter. The frame sizes below were chosen as representative of the packet sizes seen in operational networks.

- 64 bytes
- 82 bytes (minimum frame size for IPv6)
- 512 bytes
- 1500 bytes

Table D-21 lists the DUTs and test equipment used during this test.

Table D-21 Equipment Configuration

Equipment Name		Model Number	IOS/OS Version(s)
DUT	Cisco Router	7600	12.2 (33) SRA1
	Cisco Router	3845	12.4 (9) T1
	Cisco Catalyst	6500	12.2 (18) SXF6
	Cisco Catalyst	3750	12.2 (25) SEE2
Test Equipment	Pagent	N/A	N/A
	Agilent	N/A	N/A
	Spirent	N/A	N/A
	Solar Winds	N/A	N/A

Results

Routing

The test results indicated that running OSPF, Intermediate System to Intermediate System, and BGP routing protocols in dual-stack mode scale equally well for IPv4 and IPv6. There are some definite resource considerations when deploying dual-stack mode in a network. Overall, testing indicates a 5-15% increase in CPU utilization and a doubling of memory requirements.

Subnets

Testing indicated that IPv4 and IPv6 will scale equally well and that both protocols can coexist in the same platform. In all cases, significant strain was placed on the CPU when both IPv4 and IPv6 hosts were initially connecting. Testing showed that memory utilization increased. This increase is expected because the platform now has to store L2 to L3 mapping information for IPv4 and IPv6. A rule of thumb that can be taken from the test results is that enabling dual-stack operations will increase the memory utilization for the storage of L2 to L3 mapping information by 1.5 times. This number should serve as a guideline only.

Testing also revealed that all platforms can handle both IPv4 and IPv6 ACLs equally well. The 3800 platform was only able to process packets at approximately 45% line rate. This result was expected because the 3800 CPU is involved in the processing of all packets. When the ACLs were added, this placed additional strain on the CPU, which resulted in lower packet forwarding rates. The hardware based forwarding platforms were not impacted by the addition of ACLs to the interfaces.

Multicast

Testing indicated that enabling multicast services for IPv4 and IPv6 scales equally well for both protocols. There are some resource considerations to take into account when enabling dual-stack operations in each chassis. CPU utilization was not impacted when enabling IPv6 multicast services on the various platforms. Testing did indicate that there will be an increase in memory utilization as the chassis now has to maintain multicast state for the IPv4 and IPv6 protocols. The tests showed that the IPv6 multicast information will consume approximately 1.5 times the amount of space that the IPv4 information uses. These numbers should serve as guidelines only.

IP Services

Testing indicated that enabling basic IP services HTTP and FTP for IPv4 and IPv6 scales equally well for both protocols. There are some resource considerations to take into account when enabling dual-stack operations in each Server. CPU utilization was somewhat impacted when enabling IPv6 FTP and HTTP, but the increase from IPv4-only to dual-stack IPv4/IPv6 is limited by 18% in Windows 2003 Server for HTTP and by 2% in Fedora Server for FTP. Testing indicated there will be no noticeable change in memory requirements from IPv4-only to dual-stack IPv4/IPv6 for both HTTP and FTP services under both platforms.

Conclusions

Enabling IPv6 on the various platforms did require more resources such as CPU and memory. When IPv6 was enabled on the software forwarding based platforms, there was a 10-15% increase in CPU utilization. This observation did not carry over to the platforms that forwarded packets based on hardware. Enabling IPv6 on the platforms that forward packets based on hardware did not exhibit any increase in CPU utilization.

Memory usage did not depend on whether or not the platform forwarded packets in hardware or software. Memory usage showed a definite increase in the tests when IPv6 was enabled. The observed trend was that IPv6 required 1.5 to 2 times more memory than IPv4 information structures.

Overall, the results obtained showed that IPv6 will scale to the same levels as IPv4. There are more demands placed on system resources when running IPv4 and IPv6 in dual-stack mode. These added demands must be taken into account and further tested and evaluated when integrating IPv6 into the network.

Router tests conducted in this work were confined to Cisco products. To enhance the degree of confidence in the scalability of other vendors' IPv6 router products, the DoD needs to tested them as well. Also needed are tests to confirm scalability of IPv6 Directory Services; specifically Microsoft Active Directory, which is widely used in the DoD.

D.31 Milestone Objective 2 IPv6 Scenario 1 Implementation Guide & Test Parameters

Testing Organization and Publication Date

AFIWC
15 August 2006

Summary

The AFIWC IPv6 team investigated the effect IPv6 could have on an operational Air Force network in order to provide confidence that IPv6 enabled nodes will continue to function properly in an operational network. This implementation guide consists of a configuration guidance designed to ensure that the current network will continue to function at a high level when IPv6 is enabled on various nodes.

Test and Evaluation Method

Experiment

Joint Staff Operational Criteria Tested

1 (1.4, 1.4.2)
8 (8.1, 8.1.1)

Configuration

This guide's scenario 1 links two IPv4/IPv6 Enclaves via a single End Building Node router. All clients were loaded with the Air Force Standard Desktop Configuration. This is consistent with the standard Air Force desktop configurations and deployments found in today's Air Force Enterprise networks. Table D-22 lists equipment, model numbers, and IOS/OS versions used during testing.

Table D-22 Equipment Configuration

Equipment Name		Model Number	IOS/OS Version(s)
Hardware	Cisco Router	3825	12.3(14)T7
	Cisco Ethernet Switch	3750	12.2(25)SE
	Dell Power Edge Server	1850	Microsoft Windows Enterprise Server 2003 SP1
	Dell Optiplex	GX520	XP Pro SP2
	HP Laser Jet	4000N	G.05.35
	HP Laser Jet	4250 DTN	v.31.08.ff
Test Equipment	IXIA Traffic Generator	400T	3.70.24.46.SP3a
	Ethereal	N/A	N/A
	Mozilla Firefox	N/A	N/A

Results

Symantec Anti-Virus

Symantec Corporate Edition 10.2, Symantec System Center was installed on primary services servers, which were running Windows Server 2003 to act as the antivirus servers. No problems were encountered with the updating of virus signatures. It is important to note that Symantec Corporate Edition 10.2 does not support IPv6.

DHCP

DHCPv6 did not work in a dual-stack network. Windows XP and Server 2003 do not support DHCPv6 solicitations. IPv6 addresses were manually created by using the *netsh* command line tool.

Web Services

Many difficulties were experienced in getting Internet Explorer (IE) 6.0 to recognize the literal IPv6 address format. Research revealed that IE 6.0 does not support literal IPv6 address formatting as it does IPv4. However, manipulating literal hexadecimal IPv6 addresses is usually not done, so lack of this functionality is unlikely to prevent installing and using IPv6. After creating a website, it was accessed using Mozilla's FireFox 1.5.0.6 which worked well when inputting the IPv6 address in the Uniform Resource Locator (URL) link.

Recommended Tests

Table D-23 lists tests, functions of the test, and the expected result. This will assist other agencies in testing their current and future networks with dual-stack enabled.

Table D-23 Test Summaries

Test	Function of Test	Expected Result
ICMPv6 Header Information	Deny transit of IP packets with protocol 58 with use of an ACL.	The router should be able to both filter the packet and record the results within the router's log.
Tunnel Information and Control Protocol	An ACL designed to deny/filter all TCP packets with port 3874 applied on the external interface of the IPv6 enclave.	The router should be able to both filter the packet and record the results within the router's log.
Heartbeat Protocol	An ACL designed to deny/filter all UDP packets with port 3740 applied on the external interface of the IPv6 enclave.	The router should be able to both filter the packet and record the results within the router's log.
Anything in Anything Protocol	An ACL designed to deny/filter all UDP packets with port 5072 applied on the external interface of the IPv6 enclave.	The router should be able to both filter the packet and record the results within the router's log.
Anything in Anything Beta Protocol	An ACL designed to deny/filter all UDP packets with port 8374 applied on the external interface of the IPv6 enclave.	The router should be able to both filter the packet and record the results within the router's log.
IPv6 Hop-by-Hop Option	An ACL designed to deny/filter all IP packets with protocol 0 applied on the external interface of the IPv6 enclave.	The router should be able to both filter the packet and record the results within the router's log.

Conclusions

IPv6 is still limited in the services and applications it supports, so dual-stacking hosts would be the most predominant means of communications in mixed IPv4/IPv6 environments. The tests in the MO2 research show that VLAN technology can benefit Air Force networks by managing various segments within a protected enclave. The significant findings show that continued research must be performed to block various ports and protocols from potentially exploiting mission systems within the current Air Force Enterprise networks. Other services such as DHCPv6 were not supported on the Windows XP and Server 2003 systems. Web server services using Microsoft Windows Server 2003 and IE 6.0 is limited in a dual-stack environment and do not allow access to the 6bone Internet. An Application Programming Interface exists that allows IPv6 enabled sites to be viewed through the browser, but only if DNS is set up correctly. IE 6.0 cannot browse sites using literal IPv6 addresses. When IE is configured to use a proxy server, all name resolution requests for web sites are forwarded to the proxy server. Until the proxy server is IPv6-enabled, proxy-based requests for local or remote IPv6 web pages are unsuccessful.

Further findings indicate that there is still much to learn about IPv6. It will not only change the way the Air Force performs their mission requirements, but also the security behind the current configuration of Air Force networks. More and more vendors are integrating IPv6 into the standard TCP/IP stack which will change applications' and services' means of communicating. Until IPv6 is fully deployed, it is difficult to determine how it will function within the Internet, let alone in conjunction with current protocols and ports.

D.32 MO2 Security Concerns for Microsoft Windows IPv6 Protocol

Testing Organization and Publication Date

AFIOC
1 July 2006

Summary

This guide was developed to analyze and provide configuration parameters surrounding MO2 security concerns for Microsoft Windows XP with SP 2 and Windows Server 2003 family with SP 1. The research performed was based upon MO1's findings using a dual-stacked system that is specific to Windows XP SP 2 while using a Windows 2003 Server SP 1 running on an AF Enterprise network.

Test and Evaluation Method

Engineering Analysis

Joint Staff Operational Criteria Tested

1 (1.4, 1.4.1)
2 (2.2, 2.2.1)
8 (8.1, 8.1.1)

Configuration

DNS Client

In order to configure DNS for transporting AAAA records across the network, the path, or IPv6 connectivity must be ensured, or a 6in4 transition mechanism must be enabled. The AAAA record can be transported over an IPv4 packet, which only requires that the IPv6 Helper Service be started on the Windows 2003 Server. IPv6 transport can be enabled on the server by using the Windows Support Tool called dnscmd. The tools are not installed by default and are typically located on the Windows Install CD - the support tool package, suptools.msi, can be found in the Support Tool folder.

DHCPv6

Microsoft currently does not support DHCPv6 on its Windows XP and Server 2003 Family. Coincidently, Microsoft's release of Vista/Longhorn has implemented a DHCPv6 service that is capable of only negotiating service information on a native IPv6 network. The Vista/Longhorn releases are also not capable of carrying any IPv4-specific information on the DHCPv6 service.

IE

Microsoft's IE 6.0 does not support literal IPv6 addresses in URLs as referenced in RFC 2732. The only way to access IPv6-enabled Web servers is to use the Internet extensions dynamic link library, Wininet.dll. When a browser wants to download Web pages using the dynamic link library, it uses a DNS query to see if the name of the Web server in the URL returns an IPv6 address. Literal IPv6 addresses in URLs mean that the browser is compatible with the use of ":" and "." characters as delimiters. Literal IPv6 addresses in a URL must be enclosed in brackets, "[ ]", in order to negotiate a Web site via its IPv6 address format.

Internet Information Services (IIS)

Support for IPv6 over HTTP-requests is handled well within IIS 6.0. This service can run on Windows Server 2003 with the IPv6 protocol enabled. The lab utilized this service as its preferred method of hosting the SharePoint Web site. The IIS Manager does not display IPv6 addresses as it does with IPv4. IPv6 information from the user interface cannot be manipulated as well. The IP Address Restrictions feature in IIS 6.0 does not support IPv6 addresses or IPv6 prefixes. Other properties such as the ServerBindings and SecureBindings metabase parameters do not support IPv6 addresses.

Results

Vulnerabilities

<u>ICMP Flood Attack</u>

Microsoft Windows XP and Server 2003 Family is prone to DoS attacks when IPv6 traffic is enabled. The exploit floods the system with IPv6 packets and crashes the system. IPv6 enclave systems were prone to this attack even when IPv4 packets were transmitted using protocol 58. An actual service attack killed the TCP/IP stack so no communication could be occur using its allocated address. The system would also not respond to any ipconfig/release or ipconfig/renew executions.

<u>Covert Channel Tool for IPv6 Encapsulation</u>

A tool called VoodooNet has the ability to encapsulate IPv6 inside IPv4 packets and circumvent most firewalls in use today. The open source tool has not been released to the public but preliminary findings show that it can potentially bypass security applications in the Enterprise network. Files with alternate streams of data could be encapsulated and moved across normal TCP/IP enabled networks. Under New Technology File System, files can be hidden in alternate data streams. Microsoft Windows does not come with default tools for listing alternate data streams. The IPv6 header extensions field could also allow additional mechanisms for hidden files to be injected into the networks.

Secure Neighbor Discovery (SEND)

An adversary on the same segment or one that gains access to the network can essentially fake router advertisements and wreak havoc on the network by changing hop limits or advertising prefixes for some malicious websites. Most of these precautions can be isolated in the IP layer with IPSec but misconfigured keys and issues with the distribution of keys make this an administrator's nightmare. A vulnerability exists in which neighbor discovery messages can be exploited using a DoS attack, enabling the attacker to reload the system. A network administrator is advised to block all IPv6 multicast traffic from leaving or entering the enterprise.

SEND was developed for just this type of problem, but full implementation of this protocol is still not widely accepted. Microsoft does not support the SEND protocol in Windows XP, Windows Server 2003, Windows Vista, or Windows Server Longhorn.

Conclusions

The changes being made to the IPv6 stack by Microsoft further emphasize the necessity to embrace this new technology and prepare the Air Force for real world application of the protocol. Researchers continue to develop guidelines and standardized implementation policies in order to make the transition into the six-stack world a seamless one. Unfortunately with change comes much failure, but with the continued testing of various scenarios and/or configurations the process to convert to IPv6 will eventually benefit the Air Force.

Implementation recommendations and procedures may change with the release of Microsoft's new OS. However, the industry is pushing this technology and those not prepared will be left to continued mismanagement of network operations.

D.33 Milestone Objective 2 IPv6 Scenario 1 Router Configuration Guide

Testing Organization and Publication Date

AFIOC
16 August 2006

Summary

This document describes how implementation of the IPv6 protocol can connect multiple segments and devices within a LAN. It details many routers' configuration features that allow IPv6 traffic to traverse over an IPv4 infrastructure. This report also gives guidance concerning vulnerabilities, IPv6 best practices, and security.

Test and Evaluation Method

Engineering Analysis

Joint Staff Operational Criteria Tested

1 (1.4, 1.4.1, 1.5, 1.5.1)
8 (8.1, 8.1.1)

Configuration

Certain tasks should be completed before configuring IPv6 on routers:

- Identify the boundary router that will be configured to run the dual stack. The router should have at least one static, globally routable IPv4 address
- Identify the internal base router that will be modified with the necessary ACLs in order to add a fail-safe in the network in the event that the boundary router is modified or misconfigured
- Recognize current dual-stack implementation in Cisco IOS software permits an interim network management solution, which allows applications such as Trivial File Transfer Protocol, ping, Telnet, and Traceroute to run over either an IPv4 or IPv6 transport
- Use the appropriate Cisco IOS images with IPv6 support. Each specific router model will have its available IOS that supports IPv6. The recommended criteria are that an advanced IP services version be used and the router has adequate memory
- Select an IPv6 interior routing protocol that is appropriate for the network configuration
- Configure the boundary router to utilize the routing protocol most appropriate to the network. Keep tunneled traffic on a separate routing protocol to avoid tunneled networks using the same routable transport as the bases' internal networks.

Results

Vulnerabilities

IOS HTTP Authorization Vulnerability

The HTTP server local authorization makes it possible to bypass authentication and execute any command on the device. All HTTP server services on the particular router or switch can be disabled to work around this problem.

Mitigate Viruses or Worms with an IDS/IPS

A traditional virus in no way changes with IPv6, but some may experience significant barriers to propagation because of the way IPv6 addressing was developed.

Best Practices

Implement Privacy Extensions Carefully

Although privacy extensions minimize the potential for scanning attacks, it also makes it difficult to troubleshoot and trace problems on the network.

Use Non-Standard Addresses for Critical Systems

Do not use addresses with the standard ::10 or ::20 designations for hosts. Make it more difficult for adversaries to guess the address space by using something like:AEF1.

Filter Unneeded Services at the Firewall

Any services not utilized at the base level for mission critical applications should be disabled such as 6to4 or Teredo tunneling.

Host and Application Security

All systems should be patched and secured more vigorously when using IPv6 since many appliances such as firewalls and IDSs have not been deployed with IPv6 support yet.

Filter ICMPv6

All nonessential ICMP messages should be blocked at the firewall. Nevertheless, considering MO1 requirements, all ICMPv6 messages should be blocked from leaving or entering the internal base router and enclave boundary router.

Filter Spoofed Networks From Entering the Enclave or Base Network

An administrator should restrict access to traffic originating from known source addresses. All private, reserved, and invalid source addresses should be blocked at the perimeter router or base boundary device.

Conclusions

The guidance set forth in this document should assist those configuring IPv6 in a network. The recommended tasks, best practices, and vulnerabilities identified will help ensure a smooth transition to IPv6.

D.34 2006 Ethernet Switch Comparison Report

Testing Organization and Publication Date

Communications Systems Evaluation Team / Technology Integration Center (TIC)
March 2007

Summary

This report evaluates core, building, and edge Ethernet switches provided by six vendors (Alcatel Networks, Cisco Systems, Extreme Networks, Force10 Networks, Foundry Networks, and Nortel Networks) for possible use in the Installation Information Infrastructure Modernization Program (I3MP). The U.S. Army Information Systems Engineering Command (USAISEC), TIC conducted these evaluations in the TIC evaluation facility from January through December 2006. This report includes a discussion of vendor strengths and weaknesses and provides conclusions and recommendations.

Test and Evaluation Method

Demonstration

Joint Staff Operational Criteria Tested

1 (1.4, 1.4.1)
2 (2.3, 2.3.1, 2.3.2)
3 (3.1, 3.1.1, 3.2, 3.2.1, 3.3, 3.3.1)
8 (8.1, 8.1.1, 8.1.2)
9 (9.1, 9.1.1)

Configuration

The following vendors (Table D-24) were tested in a laboratory environment with the use of automated test tools (Spirent and Ixia).

Table D-24 Tested Products

Vendor	Core	Building	Edge
Alcatel			OS6850-48 OS6850-24 OS6850-P48 OS6850-P24 OS6850-48X OS6850-24X

Table D-24 Tested Products (continued)

Vendor	Core	Building	Edge
Cisco Systems	Catalyst 6509 Catalyst 6506 Catalyst 6504 Catalyst 4510R	Catalyst 4507R Catalyst 6506 Catalyst 6504	Catalyst 3750-48 Catalyst 3750-24 Catalyst 3560-48 Catalyst 3560-24 Catalyst 4507R Catalyst 2960-48 Catalyst 2960-24
Extreme Networks	BlackDiamond 12804R BlackDiamond 12804C BlackDiamond 10808		BlackDiamond 8810 Summit X450-24t
Force10	E600i E300		S50V-48
Foundry Networks	BigIron RX-8 NetIron XMR 8000 NetIron XMR 4000 BigIron RX-4		FastIron SuperX FES X424-POE FES 4802-POE FES 2402-POE
Nortel Networks	ERS 8610 ERS 8606		ERS 8606 ERS 5530 ERS 5520-48 ERS 5520-24 ERS 5510-48 ERS 5510-24

Results

The switches were evaluated in four categories: performance, system functionality, network management, and security. Results are listed below for each vendor.

Alcatel Networks

The edge switches provided by Alcatel achieved line rate throughput in each applicable test. The system functionality test demonstrated the switches could handle network traffic effectively. The IPv6 support for security and management is not currently supported by these edge switches.

Cisco Systems

Cisco switches scored the highest in the core, building, and edge switch categories. They displayed an extremely fast failover time and met almost all VoIP requirements. Throughput for the Catalyst 6500 switches for IPv6 traffic was at or near line rate in most instances. These switches were capable of making IPv6 routing decisions in hardware. The switches can be securely managed only via Secure Shell Version 2 (SSHv2) when configured solely for IPv6.

Extreme Networks

Extreme's edge switches were equal to Cisco's edge switches for the highest score in this category. Traffic patterns were stable in the system test and the QoS mechanisms shaped traffic flows as expected. The IPv6 throughput was at or near line rate for all switches. Extreme is lacking in IPv6 management capabilities and can only be managed via SSHv2.

Force10 Networks

Force10 switches exhibited IPv6 traffic rates at or near line rate with one caveat. When configured for IPv6, only half the ports per module were non-blocking. While this falls within program guidelines, this should be taken into account for future implementation of IPv6. In addition, Force10 demonstrated a robust IPv6 management capability as good as any tested.

Foundry Networks

Foundry continued to show improvement in IPv6. Foundry was the first to show any line rate IPv6 throughput at the TIC, and they continued to show the same during 2006. Network management capabilities also continued to improve. The switches support the TIC's entire SNMP agent and web access requirements. The switches also supported HTTPS and SSHv2 to secure network management traffic. Only their element manager is lacking in IPv6 support.

Nortel Networks

This was Nortel's first time showing IPv6 capabilities on their switches. Traffic was processed at or near line rate. Nortel supports some IPv6 management features, although the transmission of SNMP traps is not supported. The switches also support SSHv2 with 3DES.

Conclusions

All six vendors provided complete and sound solutions that met I3MP requirements. All vendors also continued to improve their switches' IPv6 capabilities. All the core and building switches processed IPv6 traffic in hardware. For the first time, the majority of vendors supported some type of IPv6 management traffic, but no vendor demonstrated IPSec in IPv6 environment.

D.35 Implicit Peer Enclave Prefix Discovery Protocol (IM-PEPD) High Assurance Internet Protocol Encryptor (HAIPE) Discovery White Paper

Testing Organization and Publication Date

SPAWAR
23 March 2007

Summary

This paper introduces the Implicit IM-PEPD for implementation in HAIPE devices. This will support the Navy IPv6 routing architecture and addressing scheme and DoD enterprise-wide networking transition to IPv6 by the fourth quarter of FY 2008.

Test and Evaluation Method

Engineering Analysis

Joint Staff Operational Criteria Tested

1 (1.6, 1.6.1)

Configuration

The basic requirements of IM-PEPD include:

- The Cipher Text (CT) side and the COI Plain Text (PT) side of a HAIPE must have the same network prefix. This allows the local HAIPE to automatically determine the CT prefix of a remote HAIPE without an explicit discovery protocol.
- The CT interface address of the HAIPE must be administratively pre-assigned, that is, a fixed value called "COI identifier" will be inserted into the host portion of the HAIPE CT address, and this host portion will be the same for all HAIPEs within a COI. The purpose of assigning a fixed value to the host portion of the HAIPE's CT address is to protect the PT hosts' addresses. Thus, the prefix of the COI PT would be visible on the CT core, but the host portion of the addresses would not.
- The prefix of a COI must be summarized as a single-aggregated prefix that is routable on the CT core.

Conclusions

IM-PEPD is desirable in supporting the current Navy IPv6 addressing scheme for maximum IA architecture summarization and flexibility. IM-PEPD supports prefix aggregation (in an enclave/COI). This should result in a single aggregated prefix that is routable on the GIG CT core domain. IM-PEPD is globally scalable since it is based on global prefixes; thus it is ideal for supporting maximum route summarization necessary to prevent IPv6 route explosion on Navy networks and to save significant bandwidth on low-bandwidth RF tactical links.

D.36 JUICE 2006 Test Report Verification of IPv6 Stateless Auto-configuration, Tactical Reorganization & Network Mobility (NEMO)

Testing Organization and Publication Date

CERDEC
14 December 2006

Summary

JUICE 2006 allowed users from various communities to test the interoperability of their networks with other networks currently deployed or under development. During the exercise, CERDEC and Common Hardware Systems (CHS) demonstrated the benefits of IPv6 Auto-configuration, tactical reorganization and NEMO. Detailed specifications, requirements and tests necessary to certify the use of auto-configuration and NEMO are provided in this report. The results of this effort will provide a comprehensive understanding of auto-configuration and NEMO to IPv6 implementers throughout the Army and DoD.

Test and Evaluation Method

Experiment

Joint Staff Operational Criteria Tested

10 (10.1, 10.1.1)

Configuration

JUICE allowed users from various communities to test the interoperability of their networks with other networks currently deployed or under development. JUICE provided both a tactical backbone and an experimental backbone in IPv6.

Equipment included:

- Cisco 2600 Series Router (12.4) Beta NEMO Version
- Laptops (Windows/IPv6)
- Netgear Switches
- Ethereal
- Apache Web Server.

Results

Using the IPv6 backbone, the CERDEC/CHS team demonstrated three different IPv6 features. The three experiment scenarios and their results are as follows:

Scenario 1 - Stateless Auto-configuration

This scenario demonstrated the simplification of network setup by using auto-configuration in IPv6. In this scenario, a network node consisted of a router and three hosts attached to a single point on the JUICE backbone.

The mobile router's interface 0/0 was assigned the following address through autoconfiguration: XX:XX:486:4: ----:----:----:----. The dashes represent the 64 bits of the address derived from the machine's 48 bit MAC address and a 16 bit place holder (FFFE). This is the Extended Unique Identifier 64 address. The laptop configured with the following address: XX:XX:486:5::----.

All pings returned successfully to locations within and outside of Node 1's network. The JITC web page was displayed in Firefox with no problems. All links were working and accessible on Node 1. There was no significant detriment for the time to configure for the RA interval modifications. Whether the RA interval was increased or decreased, the time to configure stayed roughly 4 seconds.

Scenario 2 – Tactical Reorganization

This scenario builds upon scenario 1 by demonstrating several nodes being configured to the JUICE backbone. This simulates the in-theater need of creating homogenous networks out of parts.

Node 1's mobile router interface 0/0 was assigned the following address through autoconfiguration: XX:XX:486:4:----:----:----:----. The Node 2's mobile router interface 0/0 was assigned the following address through autoconfiguration: XX:XX:486:4:----:----:----:----.

All pings returned successfully to locations within and outside of the Tactical Operation Center (TOC) 1's network. The JITC web page was displayed in Firefox with no problems. All links were working and accessible on Node 1 and Node 2. There was no significant detriment for the time to configure for the RA interval modifications. Whether the RA interval was increased or decreased, the time to configure stayed roughly 4 seconds.

Scenario 3 – NEMO

Using the previously demonstrated auto-configuration, scenario 3 builds on scenarios 1 and 2 by using auto-configuration in conjunction with IPv6 NEMO. As the mobile node roams, they depend on auto-configuration to obtain care-of addresses on the mobile router that allows seamless communication between the mobile node's components and the corresponding nodes. NEMO allows networks to operate in an on-the-move capacity with minimal user intervention.

Node 1's mobile router interface 0/0 was assigned the following address through autoconfiguration: XX:XX:486:4:----:----:----:----. Node 2's mobile router interface 0/0 was assigned the following address through autoconfiguration: XX:XX:486:4:----:----:----:----.

All pings returned successfully to locations within and outside of TOC 1's network. The JITC web page was displayed in Firefox with no problems. All links were working and accessible on Node 1 and Node 2. Node 2 was successfully moved from TOC 1 to TOC 2 with connectivity to the inside hosts remaining the same.

Conclusions

Scenario 1 – Stateless Auto-configuration

In demonstrating the use of stateless auto-configuration, this experiment provided simple metrics on the time to configure on both a network device and host system. Stateless auto-configuration simplifies the time and effort necessary to set up and make a network operational. Current methods of network initialization involve hard-coded static addressing and labor intensive processes. By using stateless autoconfiguration, the Army can eliminate the errors caused by manual intervention and the need to configure and distribute DHCP servers.

Scenario 2 – Tactical Reorganization

Building upon the work of scenario 1, Tactical Reorganization provided a more "real world" application of stateless auto-configuration. This scenario was an example of how stateless auto-configuration can facilitate brigade combat teams merging into a single network with minimal user intervention. Tactical reorganization by stateless auto-configuration reduces the time and potential errors caused by current methods of manual re-configuration.

Scenario 3 – NEMO

Using the previous auto-configuration, scenario 3 built upon scenarios 1 and 2 by using auto-configuration in conjunction with IPv6 NEMO. This experiment shows how easily on-the-move units can transfer from one TOC to another without having to worry about reconfiguring their networks for the new location. Communication continues at their original addresses without user or community knowledge of the change.

Recommendations

The work of the DoD/Army to transition to IPv6 is a complex ongoing effort. The experiments outlined for the JUICE 2006 exercise for network initialization and NEMO are important but only a small segment of this effort. The recommendation of Product Director CHS and CERDEC is that more laboratory and analysis work be done with these features and additional features of IPv6 to show their benefit, usability and importance to the DoD/Army transition to IPv6. It is the firm opinion of the CHS and CERDEC team that the full impact of the operational benefits of IPv6 for net-centric systems architecture is just beginning to be comprehended.

D.37 Implementing Internet Protocol Version 6 (IPv6) on an Army Installation

Testing Organization and Publication Date

USAISEC/TIC
April 2007

Summary

This paper investigates the network service areas of a typical Army post and shows what can be achieved now with IPv6 and what lags behind in achieving IPv4 parity. It describes the current state of the industry and the pieces which need to become mature before IPv6 achieves IPv4 parity.

Test and Evaluation Method

Engineering Analysis

Joint Staff Operational Criteria Tested

1 (1.1, 1.1.1, 1.5, 1.5.1)
2 (2.2, 2.2.1, 2.2.2, 2.2.3)
8 (8.1, 8.1.1)

Configuration

In order to implement an IPv6 pilot on a post, two assumptions are made.

1. Every affected device in the system will be dual stack, supporting IPv4 and IPv6. This includes the application server, client, and network backbone. There will not be any IPv6-only devices and no tunneling.
2. The application will reside entirely on-post. The client and server machines will all be on the same post and no IPv6 traffic will leave the post. This meets the MO1 guidance.

An address plan is necessary before establishing IPv6 traffic. Most IPv6 experts suggest that a post IPv6 address plan should closely reflect the current IPv4 addressing plan, to ease network management, but opportunity exists to improve the addressing scheme in IPv6. Addresses should be given out in a manner that will facilitate hierarchical routing, where prudent, and should follow Army and DoD addressing policies. Unfortunately, Army and DoD addressing policies are not complete at this time, and so a post cannot obtain permanent IPv6 address space.

Results (Issues)

Several concerns are prevalent in any implementation of IPv6; IPSec is one of the most controversial. Current guidance states that all IPv6 devices must support IPSec. Current NSA Guidance appears to indicate any IPSec device is an IA device and therefore must undergo Federal Information Processing Standard (FIPS) certification and National Information Assurance Partnership (NIAP) Common Criteria evaluation. The majority of available IPv6 devices do not support IPSec. Both the development of IPSec capabilities and the FIPS/NIAP processes are expensive and time-consuming for vendors, meaning extensive delays in getting secure products for DoD implementations.

Another issue is that upgrades are required for most servers to support the 64-bit bus speed required for Longhorn. The Network Enterprise Technology Command has proactively mandated that future server purchases must be 64-bit, but the bulk of current servers are only 32-bit.

Finally, the issue of addressing policies is not yet defined for DoD and Army. A pilot implementation could proceed with temporary IPv6 addresses, but unless an addressing plan is defined, implementers risk wasting a great deal of time and effort in renumbering and restructuring a pilot implementation when addressing plans are finalized.

Many of the delays in DoD's IPv6 implementation occur because commercial vendors do not see the pressing need to migrate to IPv6. Twenty years ago, DoD was a dominant customer in the communications industry and DoD directives were taken very seriously by industry. Today, DoD represents a relatively small market segment for most commercial vendors. To make matters worse, DoD as a whole is not investing money into IPv6 development and is only half-heartedly promoting IPv6 implementation on its networks. It is a classic Catch-22; DoD agencies do not want to invest a lot of money into IPv6 until industry starts making better products, but industry does not want to spend a lot of money developing IPv6 products until customers start buying them.

Conclusions

Implementing IPv6 on an Army post requires many more components than just IPv6-enabled core elements. Besides the switches, implementers need to be concerned with server and client OSs, network scanning and vulnerability analysis tools, addressing plans, policies, and training. Commercial development for these aspects of IPv6 are lacking, so conducting pilots at this time is very difficult. The DoD needs to continue to encourage industry to develop IPv6 products. The DITO should presently publish a mandate requiring APL usage at some future date and encouraging vendors to submit their products for DoD APL testing. Army program managers need to pressure vendors to develop IPv6 capabilities in their products and applications and pursue testing, to confirm that they will work in the Army secure dual-stack environment.

D.38 Juniper Networks Internet Protocol Version 6 Report

Testing Organization and Publication Date

JITC
June 2007

Summary

The AIPTL at Fort Huachuca, Arizona, conducted testing of Juniper routers in a dual-stack environment. The objective of this testing was to compare the performance of IPv6 with that of IPv4. Tests evaluated individual network devices.

Test and Evaluation Method

Experiment

Joint Staff Operational Criteria Tested

3 (3.1, 3.1.1, 3.2, 3.2.1)
8 (8.1, 8.1.1, 8.1.2)

Configuration

Testing was done in a dual-stack environment which will be the operational environment in DoD networks during the IPv6 transition.

The DUT loading was done using Agilent automated test equipment. Frame throughput and frame latency data was collected for each device. Multiple connections were made from the automated test device to the DUT. DUT interfaces were chosen to ensure the offered traffic would constitute 100% of system capability. Thus, each device was tested under full load conditions.

Table D-25 lists the devices and the OSs that were tested. This table includes a vendor specific sample of equipment and OSs expected to be in the DISA inventory in the 2008 time frame.

Table D-25 Equipment Configuration

DUT/Platform	Interface	Operating System	Processing Engine
Juniper T640	10 Gbps Interfaces	JUNOS 8.1	N/A
Juniper M320	10 Gbps Interfaces	JUNOS 8.1	N/A
Juniper T320	10 Gbps Interfaces	JUNOS 8.1	N/A
Juniper M40e	OC-48 POS Interfaces	JUNOS 8.1	N/A

When evaluating throughput and latency of the DUT, several IPv4/IPv6 frame ratios were used. These ratios were 100% IPv4, 100% IPv6 and the following IPv6/IPv4 ratios 10/90, 50/50, and 90/10. Ratio testing characterized a router's performance during the DoD's IPv6 transition, with the 90/10 ratio representing an early IPv6 implementation and lower IPv6 traffic than with the 50/50 and 10/90 ratios representing mid-stage and later-stage IPv6 implementations.

Results

Table D-26 reports DUT throughput and latency data. The table shows test data for individual frame sizes in each IPv4/IPv6 ratio. The throughput and latency data results for each router were combined and an average for that individual combination was listed. As an example, the values in the first results block of column one represent results of a test running 0% IPv4 frames, and 100% IPv6 frames simultaneously. Throughput values are reported in the millions of frames and latency values in microseconds. The parallel between the frame size throughput and latency values for each IP ratio illustrates equivalency for IPv4 and IPv6.

Conclusions

This test was designed to demonstrate IPv6 throughput and latency equivalent to or better than IPv4. Test data indicates that IPv4 and IPv6 performance is equivalent on all devices with IPv6, and for some frame sizes, slightly outperforming IPv4.

Table D-26 Juniper IPv4/IPv6 DUT Comparison Data

IPv4/IPv6 Ratio %	Frame Size	Combined Avg. Throughput in Millions of Frames Per Second	Combined Avg. Latency in μs
0-100	86	9	25
	128	7	27
	256	3	28
	512	2	35
	768	2	39
	1024	1	43
	1280	1	47
	1518	1	50
100-0	86	9	24
	128	7	27
	256	3	29
	512	2	35
	768	2	39
	1024	1	43
	1280	1	47
	1518	1	50
50-50	86	9	25
	128	7	28
	256	3	31
	512	2	40
	768	2	46
	1024	1	52
	1280	1	58
	1518	1	63
90-10	86	9	37
	128	7	45
	256	3	57
	512	2	90
	768	2	123
	1024	1	154
	1280	1	186
	1518	1	214
10-90	86	9	30
	128	7	35
	256	3	43
	512	2	62
	768	2	81
	1024	1	100
	1280	1	119
	1518	1	134

D.39 Beyond Addresses: IPv6 Value for the GIG

Testing Organization and Publication Date

DISA
25 August 2006

Summary

This paper provides responses to some common misconceptions about the value of IPv6 for DoD and the information industry at large. It also describes Net-Centric benefits of implementing IPv6 with the emphasis on machine-to-machine communications in two environments. The paper also reviews IPv6 standardization work in the Global Grid Forum and IETF and emerging IPv6 applications. It then affirms the strategic value of the GIG transition to IPv6.

Test and Evaluation Method

Engineering Analysis

Joint Staff Operational Criteria Tested

1 (1.1, 1.1.1)
2 (2.2, 2.2.1)
4 (4.1, 4.1.1)

Configuration

IPv6 is vital for the GIG as described in the DoD memorandum of intent. This announcement has energized IPv6 proponents in the industry and government, but many network professionals remain nonchalant. Three things stand in the way of general IPv6 acceptance:

- DoD has a large IPv4 space.
- Organizations with limited IPv4 address use Network Address Translation (NAT).
- There is no "killer application."

Results

IPv4 Space

A fundamental difference between IPv4 and IPv6 is that with IPv4, the address allocation is an asset; it provides address holders with a competitive advantage related to its size. With IPv6, in contrast, the address allocation is a commodity. Every entity will have more addresses than it could possibly use, and the competitive advantage will be in applications and new types of networks that take advantage of the IPv6 address space. Therefore, the sooner DoD starts experimenting with IPv6 applications the greater benefit it will gain from IPv6.

NAT

NAT breaks end-to-end security mechanisms such as IPSec. Specifically, IPSec AH fails because the IP address provided by the NAT box cannot be authenticated and IPSec ESP does not allow recomputing the checksum in the encrypted TCP header. Some protocols, such as FTP, Reservation Protocol, DNS, SMTP, SIP, H.323, and SNMP, can work through NAT at the cost of additional complexity of engaging ALG. However, many of these fixes are not comprehensive, and state that "making SNMP ALGs completely transparent to all management applications is not an achievable task."

There are also architecture and implementation challenges in using NAT. The process of resolving names to addresses for hosts located in private address realms favors a client-server architecture and placement of the server in the public address realm, because the DNS/ALG cannot support secure DNS name servers in the private domain.

Killer Applications

As IPv6 gains critical mass, there will be no shortage of new applications and operational scenarios. Consider for instance an Army Colonel and a Navy Commander participating in a joint operation. In today's IPv4 world, the Colonel has an IP address from the Army's pool of addresses, the Commander's address is from the Navy's pool, and when they exchange information, packets traverse numerous routers, firewalls and other systems, possibly spanning the globe back and forth. With IPv6, every joint operation could have its own address block enabling the Colonel and Commander to communicate directly.

Conclusions

Focus on the size of the IPv6 address space frequently degenerates into a quantitative comparison between the IPv4 address allocation and the needs of existing applications. This approach conceals tremendous qualitative advantages provided by the commodity nature of the IPv6 space. The abundance and easy configuration of IPv6 addresses will enable DoD to provide its users with multiple addresses associated with various COIs and manage these addresses as users move around or change COIs. A hierarchical address assignment will result in an efficient route summarization and provide better control over routing table sizes than the fragmented IPv4 space. New IPv6 protocol features will streamline networking by eliminating IP header options and packet fragmentation and improving network functions such as address configuration, multicast, ICMP, DHCP, DNS, etc. Unique Local Addresses will serve closed GIG networks to preclude external routing and, in particular, may be considered for address assignment within the GIG Black Core. Furthermore, entirely new applications will emerge by capitalizing on IPv6-specific features such as Flow Label, extension headers and anycast.

D.40 Testing Known Vulnerabilities Against Internet Protocol Version 6 (IPv6)

Testing Organization and Publication Date

NSA
21 May 2007

Summary

The general purpose of this analysis is to identify IPv6 weaknesses by testing a collection of known vulnerabilities against the IPv6 protocol stack. This final report exposes and reports weaknesses within the IPv6 stack in an operational state and similar environments. This document is intended to provide information regarding the discovery, validation, analysis and observations of applying known vulnerabilities and attacks against IPv6. This document does not contain bias, recommendations, or potential mitigations.

Test and Evaluation Method

Demonstration

Joint Staff Operational Criteria Tested

1 (1.1, 1.1.1, 1.4, 1.4.1, 1.4.2)
8 (8.1, 8.1.1, 8.1.3)

Configuration

The test laboratory consisted of both hardware and software components configured with an OS able to handle the IPv6 stack in some capacity. The test network was divided into the management network and the victim network.

The management network consisted of a Windows XP machine and a network attack simulation device known as the Spirent ThreatEx 2700. The Spirent ThreatEx 2700 device was used to generate attacks against the target machines by injecting hostile traffic into the victim network. The Windows XP system served a dual purpose and was used to both manage the ThreatEx device via its proprietary management software and generate vulnerabilities via the ThreatEx generation software for the development of customized attacks. The devices on the management network were physically connected to each other with a Cisco 3560 switch. During each test, this network was used to push attacks to the ThreatEx 2700's management port for execution on the victim network.

The victim network consisted of a machine configured with the OS to be tested, a Fedora 5 machine with Ethereal to capture packets during testing, and a Windows XP machine for "ping" and "telnet" verification. At times, a Fedora 5 machine was used to carry out attacks that could not be performed with the ThreatEx device. The devices on the victim network were physically connected to each other with a 10/100 hub to ensure that all packets were passed to the intended victim. During each test the ThreatEx's attack port was connected to the victim network to inject the anomalous traffic. At no time did the vulnerability cross any routers or firewalls while destined for the OS under attack.

All OSs were installed on Dell Precision 450 computers and were used to host the OSs on both the victim and management networks. Each had a 3.0 Gigahertz processor, 80 GigaByte (GB) hard drive, and either 1.5 GB or 2.5 GB of Random Access Memory.

Results

Table D-27 lists the vulnerabilities, test number, threat type, OS under test, and the result of each test.

Table D-27 Results

Vulnerability Description	Test #	Threat Type	Windows XP (a)	Windows Vista (c)	Windows Vista Native Stack (f)	Windows 2003 Server (b)	Linux (d)	Solaris (e)
Operating System Specific								
IPv6 SYN Vulnerability	10	DoS	Fail	Fail	Fail	Fail	Fail	Fail
IPv6 avoidance of Microsoft Firewall	23	Covert Channel	Fail	Pass	Pass	Fail	N/A	N/A
Protocol Specific								
IPv6 Land Attack	1	DoS	Fail	Pass	Pass	Fail	Pass	Pass
IPv6 SYN Flood	2	DoS	Fail	Fail	Fail	Pass	Fail	Fail
IPv6 SYN Flood 2	3	DoS	Pass	Pass	Pass	Pass	Pass	Fail
RST Flood IPv6	4	DoS	Pass	Pass	Pass	Pass	Pass	Pass
FIN Flood IPv6	5	DoS	Fail	Fail	Fail	Pass	Pass	Pass
ACK Flood IPv6	6	DoS	Fail	Fail	Fail	Pass	Fail	Fail
URG Flood IPv6	7	DoS	Fail	Fail	Fail	Pass	Pass	Pass
CWR Flood IPv6	8	DoS	Fail	Fail	Fail	Pass	Pass	Pass
ECE Flood IPv6	9	DoS	Fail	Fail	Fail	Pass	Pass	Pass
IPV6 Ping flood	11	DoS	Fail	Fail	Fail	Pass	Fail	Fail
Zero Length IPv6 Packet	12	DoS	Fail	Pass	Pass	Pass	Pass	Pass
Link Local Scope Multicast Smurf attack	13	DoS	Pass	Pass	Pass	Pass	Pass	Fail
Interface Local Scope Multicast Smurf attack	14	DoS	Fail	Pass	Pass	Pass	Pass	Pass
Link Local Scope Multicast Reverse Smurf attack	15	DoS	Pass	Pass	Pass	Pass	Pass	Fail
IPv6 Random Length Field	16	DoS	Pass	Pass	Pass	Pass	Pass	Pass
IPv6 Random Length Field	16	Remote Access	Pass	Pass	Pass	Pass	Pass	Pass
IPv6 Random Priority and Flow Labels	17	DoS	Fail	Pass	Pass	Pass	Pass	Pass
Multicast done DoS	18	DoS	Pass	Pass	Pass	Pass	Pass	Pass
Interface Local Multicast ping Spoof Response	19	RFC 2460	Pass	Pass	Pass	Pass	Pass	Pass
Link Local Multicast ping Spoof Response	20	RFC 2460	Pass	Pass	Pass	Pass	Fail	Fail
IPv6 Network Device Discovery DoS (Startup)	21	DoS	Pass	Pass	Pass	Pass	Fail	Pass
IPv6 Network Device Discovery DoS (Reinitialize IPv6)	22	DoS	Fail	Fail	Fail	Fail	Fail	Fail

Table D-27 Results (continued)

Vulnerability Description	Test #	Threat Type	Windows XP (a)	Windows Vista (c)	Windows Vista Native Stack (f)	Windows 2003 Server (b)	Linux (d)	Solaris (e)
IPv6 First Fragment Flood	24	DoS	Pass	Pass	Pass	Pass	Pass	Pass
IPv6 Fragment Flood	25	DoS	Pass	Pass	Pass	Pass	Pass	Pass
Link Local Spoofing	26	Session Hijacking	Fail	Fail	Fail	Fail	Fail	Fail
IPv6 ICMP_ECHO Request Payload	27	Covert Channel	Fail	Fail	Fail	Fail	Fail	Fail
IPv6 ICMP_ECHO Request ID Header Field	28	Covert Channel	Fail	Fail	Fail	Fail	Fail	Fail
IPv6 Remote DoS - 6to4 Tunneled Smurf	29	DoS	Pass	Pass	Pass	Pass	Pass	Pass
6 to 4 Tunneled Ping Random IP	47	Covert Channel	Pass	Pass	Pass	Pass	Fail	Fail
6 to 4 Tunneled Ping Link Local IP	48	Covert Channel	Pass	Pass	Pass	Pass	Fail	Fail
6to4 Tunneled IPv6 ICMP_ECHO Request Payload	49	Covert Channel	Pass	Pass	Pass	Pass	Fail	Pass
6to4 Tunneled IPv6 ICMP_ECHO Request ID Header Field	50	Covert Channel	Pass	Pass	Pass	Pass	Pass	Pass
Fuzzing								
Fuzz all Fields	30	FUZZ	Fail	Pass	Pass	Fail	Pass	Fail
Fuzz Dest IP	31	FUZZ	Pass	Pass	Pass	Pass	Pass	Pass
Fuzz Dest MAC	32	FUZZ	Pass	Pass	Pass	Pass	Pass	Pass
Fuzz Dest Multicast	33	FUZZ	Pass	Pass	Pass	Pass	Pass	Pass
Fuzz Flow Label	34	FUZZ	Pass	Pass	Pass	Pass	Pass	Pass
Fuzz Hop Limit	35	FUZZ	Pass	Pass	Pass	Pass	Pass	Pass
Fuzz Next Header	36	FUZZ	Fail	Pass	Pass	Fail	Pass	Fail
Fuzz Payload Length	37	FUZZ	Pass	Pass	Pass	Pass	Pass	Pass
Fuzz Priority	38	FUZZ	Pass	Pass	Pass	Pass	Pass	Pass
Fuzz Source IP	39	FUZZ	Pass	Pass	Pass	Pass	Pass	Pass
Fuzz Source MAC	40	FUZZ	Pass	Pass	Pass	Pass	Pass	Pass
Fuzz Payload Length2	41	FUZZ	Fail	Pass	Pass	Fail	Pass	Fail
Fuzz Destination Mulitcast2	42	FUZZ	Pass	Pass	Pass	Pass	Pass	Pass
Fuzz Destination Mulitcast3	43	FUZZ	Pass	Pass	Pass	Pass	Pass	Pass
Fuzz Dest IP - Link Local	44	FUZZ	Pass	Pass	Pass	Pass	Pass	Pass
Fuzz Source IP - Link Local	45	FUZZ	Pass	Pass	Pass	Pass	Pass	Pass
Fuzz Version	46	FUZZ	Pass	Pass	Pass	Pass	Pass	Pass

Conclusions/Summary

Five different OSs and their TCP/IPv6 stack implementations were tested. All OSs were tested under the same conditions. Below is a short summary of each OS.

Windows XP

During DoS attacks within the link local network, Windows XP accepted TCP/IPv6 packets whose source IPv6 addresses were random and consistent with non link local source addresses. Other OS implementations discarded these packets, but Windows XP accepted and processed them.

Windows Vista – Dual Stack

During DoS attacks within the link local network, Windows Vista accepted TCP/IPv6 packets whose source IPv6 addresses were inconsistent with link local source addresses. Other OS implementations discarded these packets, but Windows Vista accepted and processed them.

Windows Vista – Native Stack

Test results were the same for the Windows Vista native stack as they were for dual stack, except for the DAD DoS test. The Windows Vista native IPv6 implementation failed the DAD test during system reboot and start-up, which differed from dual-stack testing.

Windows 2003 Server

Windows 2003 Server seems to be very resilient to most of the DoS attacks. Many of these attacks were particular floods that had some effect but none degraded overall performance. One of the only DoS attacks that was effective during testing was the Land attack when it took approximately 20 seconds after the last packet was sent to recover to a steady state.

Linux

The Fedora Core 5 OS was also tested for possible flaws in the tunneling of IPv6 packets inside of IPv4 packets. During the testing, Fedora responded to some of the tests, which were an ICMPv6 packet inside an IPv4 packet with both random external and link local source addresses. These packets solicited responses in the form of Protocol Unreachable packets. This behavior should be limited due to the possibility it may be used for enumeration or reconnaissance.

Solaris

Solaris 9 was the only OS that failed the TCP/IPv6 SYN flood attack generated from the same source IPv6 address and port. SYN2 was a SYN flood attack that did not generate random sources for each packet. Both routable and link local IPv6 source addresses from the attack packets generated NS responses from the Solaris 9 IPv6 stack implementation. It appears that these responses to the same course address caused DoS.

D.41 IPv6 Mitigation Planning Phase 3: Custom Configuration Guidance

Testing Organization and Publication Date

NSA
1 June 2007

Summary

This document addresses mitigations for the highest priority vulnerabilities that were validated during FY 2006 efforts and is the second of five mitigation reports. The purpose of this mitigation plan (Phase 3) is to implement custom configuration guidance on each OS then to re-test the failed attacks from the FY 2006 efforts as a means to identify weaknesses in the technical security controls inherent within IPv6.

Test and Evaluation Method

Experiment

Joint Staff Operational Criteria Tested

1 (1.1, 1.1.1, 1.4, 1.4.1, 1.4.2)
8 (8.1, 8.1.1, 8.1.3)

Configuration

The IPv6 test team tested five different OSs and their TCP/IPv6 stack implementations. Each system applied vendor patches, implemented Secure Technical Implementation Guides (STIGs), and applied custom configuration guidance for each test. The STIGs used to guide OS testing required that each OS be updated to the latest vendor patch level. As a result, each system was updated with the recommended patches as of November 2006. For each OS, the integrated patch utility was used to load all suggested vendor patches.

Results

This is the fourth iteration of testing, which is a Phase 3 of FY 2007 effort. The matrix listed as Table X contains five types of boxes, which denote five types of test results. First, all empty boxes are tests that passed the first round of testing (FY06 Testing Known Vulnerabilities against IPv6) and were deemed not vulnerable. Second, all green boxes represent tests that were mitigated by applying Vendor Patches [(VP) used in Table D-28; Phase 1]. Third, all blue boxes represent tests that were mitigated by the implementation of the DISA STIGs (Phase 2). Fourth, the yellow boxes represent tests that were mitigated by the implementation of Custom Configuration Guidance (CCG used in Table D-28; Phase 3). Fifth, the red boxes represent tests that failed during FY 2006 testing, vendor patch, STIG and custom configuration guidance testing. The tests denoted with red boxes are considered RFC, protocol or vendor implementation issues (Phase 4).

Table D-28 IPv6 Testing Matrix

Vulnerability Description	Test #	Threat Type	Windows XP (a)	Windows Vista (c)	Windows Vista Native Stack (f)	Windows 2003 Server (b)	Linux (d)	Solaris (e)
Operating System Specific								
IPv6 SYN Vulnerability	10	DoS	VP	Not Mitigated	Not Mitigated	VP	VP	CCG
IPv6 avoidance of Microsoft Firewall	23	Covert Channel	VP			VP		
Protocol Specific								
IPv6 Land Attack	1	DoS	VP			VP		
IPv6 SYN Flood	2	DoS	VP	Not Mitigated	Not Mitigated		STIG	CCG
IPv6 SYN Flood 2	3	DoS						VP
RST Flood IPv6	4	DoS						
FIN Flood IPv6	5	DoS	VP	CCG	CCG			
ACK Flood IPv6	6	DoS	VP	CCG	CCG		STIG	Not Mitigated
URG Flood IPv6	7	DoS	VP	CCG	CCG			
CWR Flood IPv6	8	DoS	VP	CCG	CCG			
ECE Flood IPv6	9	DoS	VP	CCG	CCG			
IPV6 Ping flood	11	DoS	VP	Not Mitigated	Not Mitigated		CCG	CCG
Zero Length IPv6 Packet	12	DoS	VP					
Link Local Scope Multicast Smurf attack	13	DoS						VP
Interface Local Scope Multicast Smurf attack	14	DoS	STIG					
Link Local Scope Multicast Reverse Smurf attack	15	DoS						Not Mitigated
IPv6 Random Length Field	16	DoS						
IPv6 Random Length Field	16	Remote Access						
IPv6 Random Priority and Flow Labels	17	DoS	VP					
Multicast done DoS	18	DoS						
Interface Local Multicast ping Spoof Response	19	RFC 2460						
Link Local Multicast ping Spoof Response	20	RFC 2460					CCG	CCG
IPv6 Network Device Discovery DoS (Startup)	21	DoS			Not Mitigated		CCG	
IPv6 Network Device Discovery DoS (Reinitialize IPv6)	22	DoS	Not Mitigated	Not Mitigated	Not Mitigated	Not Mitigated	CCG	CCG
IPv6 First Fragment Flood	24	DoS						
IPv6 Fragment Flood	25	DoS						
Link Local Spoofing	26	Session Hijacking	VP	CCG	CCG	CCG	STIG	STIG
IPv6 ICMP_ECHO Request Payload	27	Covert Channel	VP	CCG	CCG	CCG	CCG	STIG
IPv6 ICMP_ECHO Request ID Header Field	28	Covert Channel	VP	CCG	CCG	CCG	CCG	STIG
IPv6 Remote DoS - 6to4 Tunneled Smurf	29	DoS						
6 to 4 Tunneled Ping Random IP	47	Covert Channel					CCG	CCG
6 to 4 Tunneled Ping Link Local IP	48	Covert Channel					CCG	CCG

Table D-28 IPv6 Testing Matrix (continued)

Vulnerability Description	Test #	Threat Type	Windows XP (a)	Windows Vista (c)	Windows Vista Native Stack (f)	Windows 2003 Server (b)	Linux (d)	Solaris (e)
6to4 Tunneled IPv6 ICMP_ECHO Request Payload	49	Covert Channel					CCG	
6to4 Tunneled IPv6 ICMP_ECHO Request ID Header Field	50	Covert Channel						
Fuzzing								
Fuzz all Fields	30	FUZZ	VP			CCG		STIG
Fuzz Dest IP	31	FUZZ						
Fuzz Dest MAC	32	FUZZ						
Fuzz Dest Multicast	33	FUZZ						
Fuzz Flow Label	34	FUZZ						
Fuzz Hop Limit	35	FUZZ						
Fuzz Next Header	36	FUZZ	VP			CCG		STIG
Fuzz Payload Length	37	FUZZ						
Fuzz Priority	38	FUZZ						
Fuzz Source IP	39	FUZZ						
Fuzz Source MAC	40	FUZZ						
Fuzz Payload Length2	41	FUZZ	VP			CCG		STIG
Fuzz Destination Mulitcast2	42	FUZZ						
Fuzz Destination Mulitcast3	43	FUZZ						
Fuzz Dest IP - Link Local	44	FUZZ						
Fuzz Source IP - Link Local	45	FUZZ						
Fuzz Version	46	FUZZ						

Conclusions/Summary

The IPv6 test team tested five different OSs and their TCP/IPv6 stack implementations. During this phase of mitigation testing, 85% of the vulnerabilities passed testing after installation of vendor patches, implementation of STIGs and custom configurations. Of the 50 tests conducted (or the remaining failures from the STIG implementation testing) 37 were mitigated and recognized as “Pass”.

D.42 IPv6 Mitigation Planning Phase 4: RFCs and Protocols

Testing Organization and Publication Date

NSA
1 June 2007

Summary

This document presents the results of the remaining failures after vendor patches, STIGs, and custom configuration guidance had been applied. This document identifies the remaining failures that are either RFCs, or vendor or standards bodies' implementation issues and designates them as a deferred risk (Phase 4). This mitigation report is the last report intended as a summarization of the effort to expose and report weaknesses within the implementation of IPv6 in an operational state and similar environments once OS patches, STIGs and custom configuration guidance have been applied.

Test and Evaluation Method

Engineering Analysis

Joint Staff Operational Criteria Tested

1 (1.1, 1.1.1, 1.4, 1.4.1, 1.4.2)
8 (8.1, 8.1.1, 8.1.3)

Configuration

The failures that were deemed RFC, vendor implementation or standards body issues became known as deferred risks. The team completed research on each of the remaining failures to determine whether the remaining failure was a vendor or RFC issue and what current steps were being taken to mitigate the issue, if any.

Results

Windows XP

During the custom configuration testing of Windows XP, 1 vulnerability was tested. After exhausting possibilities through a custom configuration, Duplicate Address Discovery DoS (reinitialized) still failed to be mitigated within Windows XP.

Windows 2003 Server

During the custom configuration testing of Windows 2003, a known vulnerability was tested. After exhausting possibilities for mitigations through a custom configuration, one vulnerability

(Duplicate Address Discovery DoS reinitialized) failed to be mitigated within Windows Server 2003.

Linux

During the custom configuration testing of Fedora Core 5, 8 known vulnerabilities were tested. After testing was completed, all vulnerabilities had been mitigated.

Solaris

During the custom configuration testing of Solaris 9, 9 known vulnerabilities were tested. After exhausting possibilities for mitigations through a custom configuration, 2 vulnerabilities (ACK flood IPv6 attack and Link Local Scope Multicast Smurf attack) failed to be mitigated within Solaris 9.

Windows Vista Dual Stack

During the custom configuration testing of Windows Vista Dual Stack configuration, 12 known vulnerabilities were tested. After exhausting possibilities for mitigations through a custom configuration, 4 vulnerabilities (IPv6 SYN Flood, IPv6 Ping Flood, Duplicate Address Discovery DoS reinitialized, and IPv6 SYN Flood) failed to be mitigated within Windows Vista Dual Stack configuration.

Windows Vista Native Stack

During the custom configuration testing of Windows Vista Native IPv6 Stack configuration, 13 known vulnerabilities were tested. After exhausting possibilities for mitigations through a custom configuration, 5 vulnerabilities (IPv6 SYN Flood, IPv6 Ping Flood, Duplicate Address Discovery DoS reinitialized, Duplicate Address Discovery DoS Startup, and IPv6 SYN Flood) failed to be mitigated within Windows Vista Native IPv6 Stack.

Conclusions/Summary

After the completion of vendor patch, STIG and custom configuration guidance implementations, tests that passed were deemed mitigated; those that did not pass were deemed failures or vulnerabilities. Vendor patch, STIG, and custom configuration guidance testing yielded 13 remaining failures that were not mitigated by hardened systems. The remaining 13 failures are vendor implementation or RFC issues and are mitigated through deferring the risk.

D.43 Internet Protocol Version 6 (IPv6) Mitigation Plan Phase 1: Vendor Patch Implementation Plan

Testing Organization and Publication Date

NSA
21 May 2007

Summary

This document addresses mitigations for the highest priority vulnerabilities that were validated during FY 2006 efforts and is the first of five mitigation reports. The general purpose of this mitigation plan (Phase 1) is to implement vendor patches on each OS and re-test the failed attacks from the FY 2006 effort, as a means to identify weaknesses in the technical security controls inherent within IPv6.

Test and Evaluation Method

Experiment

Joint Staff Operational Criteria Tested

1 (1.1, 1.1.1, 1.4, 1.4.1, 1.4.2)
8 (8.1, 8.1.1, 8.1.3)

Configuration

The IPv6 test team analyzed five different OSs and their TCP/IPv6 stack implementations. All OSs were updated to the latest vendor patch levels as of November 2006. For each OS, the built in patch utility was used to load every possible patch that was suggested. This ensured that all patches provided by the vendors were included.

Results

In FY 2006, efforts yielded 84 failed attacks or tests across five OSs. After vendor patches were applied, 60 failed attacks remained. Failed attacks were reduced by 28%. Table D-28 (pages 138, 139), shows what attacks/tests were corrected by vendor patches. Table D-29 gives an overall view of which OS was tested, the number of failed tests, and number of failed tests after vendor patches were applied.

Table D-29 Failed Attacks After Implementing Vendor Patches

Operating Systems Tested	Number of Failed Tests	Number of Failed Tests After Vendor Patches are Applied
Windows XP	20	2 (10%)
Windows 2003 Server	10	7 (70%)
Windows Vista Native	13	13 (100%)
Windows Vista Dual Stack	12	12 (100%)
Linux – Fedora Core 5	12	11 (92%)
Unix – Solaris 9	17	15 (88%)
Total Attacks	**84**	**60 (72%)**

Conclusions/Summary

The IPv6 Testing investigated five different OSs and their TCP/IPv6 stack implementations. All OSs were tested under the same conditions and results were reported accordingly. During this phase of mitigation testing, 29% of the vulnerabilities passed testing by installation of vendor patches. Of the 85 tests, 24 passed.

D.44 Internet Protocol Version 6 (IPv6) Mitigation Plan Phase 2: STIG Implementation

Testing Organization and Publication Date

NSA
21 May 2007

Summary

This document addresses mitigations for the highest priority vulnerabilities that were validated during FY 2006 efforts and is the second of five mitigation reports. The purpose of this mitigation plan (Phase 2) is to implement OS STIGs for each OS, then to re-test the failed attacks from the FY 2006 efforts as a means to identify weaknesses in the technical security controls inherent in IPv6.

Test and Evaluation Method

Experiment

Experiment

Joint Staff Operational Criteria Tested

1 (1.1, 1.1.1, 1.4, 1.4.1, 1.4.2)
8 (8.1, 8.1.1, 8.1.3)

Configuration

The IPv6 test team assessed five different OSs and their TCP/IPv6 stack implementations. The STIGs used to guide OS testing required that each OS be updated to the latest vendor patch level. As a result, each system was updated with the recommended patches, as of November 2006. For each OS, the integrated patch utility was used to load all suggested vendor patches.

Results

The IPv6 team applied vendor patches then ran 84 tests. Sixty tests remained prior to DISA STIG implementations. Of these 60 tests, 10 tests passed or were mitigated by STIG implementation. Based on these results, vendor patches and STIGs mitigated 40% of the vulnerabilities discovered in the FY 2006 IPv6 Testing. Table D-28 (pages 138, 139), shows which attacks/tests were mitigated by STIGs. Table D-30 gives an overall view of which OS was tested, the number of failed tests, the number of failed tests after vendor patches were applied, and the number of failed tests after STIGs and vendor patches were applied.

Table D-30 Failed Attacks Post STIG Implementation

Operating Systems Tested	Number of Failed Tests (FY06 Testing)	Number of Failed Tests After Vendor Patches are Applied	Number of Failed Tests After STIGs and Vendor Patches are Applied
Windows XP	20	2 (10%)	1 (5%)
Windows 2003 Server	10	7 (70%)	7 (70%)
Windows Vista Native	13	13 (100%)	13 (100%)
Windows Vista Dual Stack	12	12 (100%)	12 (100%)
Linux – Fedora Core 5	12	11 (92%)	8 (66%)
Unix – Solaris 9	17	15 (88%)	9 (53%)
Total Attacks	**84**	**60 (72%)**	**50 (60%)**

Conclusions/Summary

The IPv6 test team analyzed five different OSs and their TCP/IPv6 stack implementations. All OSs were tested under the same conditions, and results were reported accordingly. During this phase of mitigation testing, 40% of the original failures had passed testing after installation of vendor patches and implementation of STIGs (50 failures remained out of the original 84 failures identified during the initial vulnerability testing). Of the 60 tests conducted (or the remaining failures from the vendor patch implementation testing), 10 were recognized as Pass.